

**EXHIBIT 148 (PART 1) TO  
HARVEY DECLARATION  
REDACTED VERSION**

WILMERHALE

Leon B. Greenfield

+1 202 663 6972 (t)  
+1 202 663 6363 (f)  
leon.greenfield@wilmerhale.com

January 25, 2010

Ryan S. Struve, Esq.  
Networks and Technology Enforcement Section  
U.S. Department of Justice  
Antitrust Division  
450 5th Street, NW, Suite 7100  
Washington, DC 20530

Re: CID No. 25507, Issued to Intel Corporation

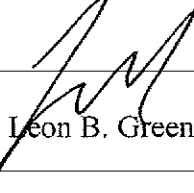
Dear Ryan:

Enclosed please find a submission on behalf of Intel Corporation in connection with the above-captioned matter.

The paper and supporting documents submitted today contain proprietary and highly confidential information. Intel requests that they be treated confidentially and not disclosed to anyone outside the Department of Justice. We expect that the Department will provide the materials with all the protections from disclosure under 15 U.S.C. § 18a(h) and all other applicable laws and rules. If this is not correct, please inform me at your earliest convenience.

Please contact me or Stacy Humes-Schulz (202-663-6076) if you have any questions.

Sincerely,



Leon B. Greenfield

Enclosure

**From:** Humes-Schulz, Stacy  
**To:** 'ryan.struve@usdoj.gov'  
**CC:** Greenfield, Leon  
**Sent:** 1/25/2010 4:53:21 PM  
**Subject:** Intel White Paper  
**Attachments:** INTEL WHITE PAPER 1-25-10.pdf; Intel White Paper Cover Letter.PDF; Intel White Paper Exhibits 1-37 (2).pdf.PDF

Ryan,

Attached please find an electronic copy of Intel's white paper and supporting exhibits. Hard copies will be sent over by courier this afternoon.

Thanks,

Stacy

**Stacy Humes-Schulz | WilmerHale**

1875 Pennsylvania Avenue NW

Washington, DC 20006 USA

+1 202 663 6076 (t)

+1 202 663 6363 (f)

[stacy.humes-schulz@wilmerhale.com](mailto:stacy.humes-schulz@wilmerhale.com)

**Please consider the environment before printing this email.**

This email message and any attachments are being sent by Wilmer Cutler Pickering Hale and Dorr LLP, are confidential, and may be privileged. If you are not the intended recipient, please notify us immediately—by replying to this message or by sending an email to [postmaster@wilmerhale.com](mailto:postmaster@wilmerhale.com)—and destroy all copies of this message and any attachments. Thank you.

For more information about WilmerHale, please visit us at <http://www.wilmerhale.com>.

**WILMERHALE**

**Leon B. Greenfield**

+1 202 663 6972(t)

+1 202 663 6363(f)

leon.greenfield@wilmerhale.com

January 25, 2010

Ryan S. Struve, Esq.  
Networks and Technology Enforcement Section  
U.S. Department of Justice  
Antitrust Division  
450 5th Street, NW, Suite 7100  
Washington, DC 20530

Re: CID No. 25507. Issued to Intel Corporation

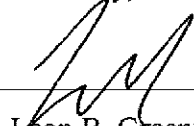
Dear Ryan:

Enclosed please find a submission on behalf of Intel Corporation in connection with the above-captioned matter.

The paper and supporting documents submitted today contain proprietary and highly confidential information. Intel requests that they be treated confidentially and not disclosed to anyone outside the Department of Justice. We expect that the Department will provide the materials with all the protections from disclosure under 15 U.S.C. § 18a(h) and all other applicable laws and rules. If this is not correct, please inform me at your earliest convenience.

Please contact me or Stacy Humes-Schulz (202-663-6076) if you have any questions.

Sincerely,



Leon B. Greenfield

Enclosure

CONTAINS HIGHLY CONFIDENTIAL BUSINESS INFORMATION  
PROTECTED FROM DISCLOSURE UNDER  
15 U.S.C. § 57b-2 AND 16 C.F.R. §§ 4.10-4.11

**SUBMISSION OF INTEL CORPORATION  
REGARDING DEPARTMENT OF JUSTICE CID**

January 25, 2010

Leon B. Greenfield  
Avery W. Gardiner  
Stacy Humes-Schulz  
Wilmer Cutler Pickering Hale & Dorr, LLP  
1875 Pennsylvania Avenue, NW  
Washington, DC 20006

We understand that the Antitrust Division Staff is considering whether to recommend that the Department of Justice bring a per se Section 1 case against Intel based on alleged “agreements”<sup>1</sup> with Pixar, Dell, and Google to eliminate one method of recruiting each other’s employees – uninvited cold calling. The alleged agreements were plainly not per se illegal and easily withstand rule of reason scrutiny. Indeed, it is implausible that the alleged agreements actually injured competition in any employment market, and we do not understand the Staff to contend otherwise. Intel respectfully submits that the Department should close its investigation without taking action.

The alleged agreements must be evaluated under the rule of reason because they are ancillary restraints that are reasonably necessary to efficiency-enhancing collaborations between the parties. *See* 2000 Department of Justice and Federal Trade Commission Antitrust Guidelines for Collaborations Among Competitors § 3.2. Moreover, the rule of reason would apply even if the alleged agreements were not deemed ancillary restraints. Per se treatment is reserved only for business practices that are “so plainly anticompetitive that no elaborate study of the industry is needed to establish their illegality.” *Texaco Inc. v. Dagher*, 547 U.S. 1, 5 (2006) (quoting *Nat’l Soc’y of Prof’l Eng’rs v. United States*, 435 U.S. 679, 692 (1978)). Conduct is branded unlawful per se only “after considerable experience with [those] business relationships . . .” demonstrates that the conduct is always or nearly always anticompetitive. *United States v. Topco Assocs.*, 405 U.S. 596, 607 (1972). Those narrow conditions do not apply here.

<sup>1</sup> We use the terms “agreed” and “agreement” here only to simplify the discussion. By doing so, we do not concede that Intel’s recruiting practices constitute “agreements” for Sherman Act § 1 purposes.

CONFIDENTIAL BUSINESS INFORMATION

The Pixar, Dell, and Google alleged agreements are lawful under the rule of reason. They are – at most – bilateral agreements that facilitated pro-competitive collaborations and had very limited, if any, impact on the parties’ recruiting of employees – even between each other, let alone any anticompetitive effect on employment markets. None of the alleged agreements kept the parties from hiring each other’s employees or from actively soliciting employees who indicated in any way they were open to considering new employment – *e.g.*, by submitting a resume on Intel.com/jobs, approaching Intel at a jobs fair, or networking with Intel employees.<sup>2</sup> Moreover, restrictions on cold calling had little impact in the context of Intel’s hiring practices, given that Intel only sparingly employed this inefficient and largely unproductive recruiting method. Indeed, as is explained in section IV(D), *infra*, less than two percent of the candidates Intel hired in 2009 came to Intel through staffing department cold calls.

The Staff’s attempt to apply the per se rule to these facts would turn the rule on its head. The purpose of the per se rule is to avoid the need for detailed inquiry into competitive effects *when the practice at issue so obviously harms competition that no such inquiry is necessary*. *See, e.g., Dagher*, 447 U.S. at 5. Here, however, the Division would be pursuing a per se theory not because the conduct at issue would invariably injure competition, but rather to condemn pro-competitive agreements that *could never have injured competition* in a relevant market.

For these reasons, we respectfully submit that the Antitrust Division should close its investigation of Intel without bringing any action.

\* \* \* \* \*

<sup>2</sup> As the Staff knows, Intel’s staffing organization and Legal Department developed a “Courtesy List,” a compilation of legal advice to Intel’s staffing organization regarding certain companies (including Pixar) for which Intel has special recruiting protocols in place. Intel produced the Courtesy List to the Division on November 23, 2009. This was not a no-hire or even no-solicit agreement but simply unilateral guidelines for Intel’s staffing organization.

CONFIDENTIAL BUSINESS INFORMATION

This submission first describes the facts relating to the alleged agreements involving Pixar, Dell, and Google. Next, we show that that these alleged agreements were ancillary to efficiency-creating joint collaborations, thus mandating application of the rule of reason. We then explain that the rule of reason applies even if the Department does not view these as ancillary agreements. Finally, we show that the alleged agreements are lawful under the rule of reason because they brought no anticompetitive effects in a properly defined market and advanced pro-competitive benefits.

# **I. Background Relating to the Alleged Agreements to Limit Uninvited Cold Calling**

## *A. Pixar*

### 1. Scope of Joint Collaboration

Pixar and Intel have worked together on various joint development projects since the 1990s. Initially, Intel's relationship with Pixar was a supplier-customer relationship, with Intel supplying Pixar Xeon microprocessors for servers. Over time, that relationship intensified as Intel began increasing its graphics capabilities and working directly with Pixar to optimize Pixar's graphics "rendering" software, which runs on Intel architecture.

For example, in 2002, Intel dedicated an engineer to optimizing portions of Pixar's code. Intel also collaborated with Pixar in the production of the 2007 movie "Ratatouille," both on technical work and in marketing the movie. In a contemporaneous press release, Pixar's Senior Vice President for Technology described the technical collaboration: "There were many technical challenges in the making of 'Ratatouille.' Intel's advanced computing capabilities helped Pixar bring 'Ratatouille' to life faster than ever, delivering a 30 percent performance improvement in the computer-generated animation and visual effects rendering software. Faster



rendering gives lighting designers more time to create even more realistic images like an animated Eiffel Tower or Remy the rat's winning smile."<sup>3</sup> Pixar exhibited Intel microprocessor-based workstations at movie theatres to highlight the Intel technology powering its advanced animation.

The companies' collaboration continued to grow as Intel expanded its graphics capability in the mid to late 2000s, including by acquiring Neoptica in the fall of 2007.<sup>4</sup> Neoptica, a small company focused on increasing the speed of graphics applications,<sup>5</sup> had employees with strong rendering skills, which increased Intel's ability to collaborate with Pixar. About the time it acquired Neoptica, Intel began working with Pixar on Intel's plan to enter the discrete graphic processing sector with its "Larrabee" product.<sup>6</sup> Intel asked Pixar to analyze its Larrabee plans and to help develop Larrabee as a solution to Pixar's need for increased rendering speed. Additionally, Intel and Pixar are working together to improve both companies' power-tuning capabilities, which would allow Pixar to reduce power consumption related to rendering.

These projects have involved close collaborations between Intel and Pixar, especially between the companies' engineers, who routinely work closely with each other in person or otherwise. Intel and Pixar conduct bi-weekly conference calls to discuss their various projects. Intel also has an Enterprise Account Manager dedicated to Pixar, who is responsible for educating Pixar on Intel's technology and systems and developing opportunities for the two companies to work together.

## 2. Intel's Uninvited Solicitation Policy Regarding Pixar

<sup>3</sup> Exhibit 1, <http://www.intel.com/pressroom/archive/releases/2007/20070703corp.htm>.

<sup>4</sup> Intel continued to build its graphics capability by acquiring Swiftfoot Graphics for its graphical rendering technology in the first quarter of 2009.

<sup>5</sup> Exhibit 2, [http://www.xbitlabs.com/news/vidco/display/20071126140151\\_Intel\\_Quickly\\_Acquires\\_Graphics\\_Software\\_Firm.html](http://www.xbitlabs.com/news/vidco/display/20071126140151_Intel_Quickly_Acquires_Graphics_Software_Firm.html).

<sup>6</sup> Exhibit 3, <http://software.intel.com/file/18198/>.

CONFIDENTIAL BUSINESS INFORMATION

In October 2008, Pixar's Senior Vice President for Technology Greg Brandeau complained to Pat Gelsinger, a then-Intel executive, about Intel's uninvited cold calling of Pixar employees. Pixar's complaint to Intel was tied directly to the companies' partnership: "Given that Pixar and Intel do have a close relationship, could you ask the recruiters inside of Intel to please stop recruiting our people? I am not saying the crazy thing that people at Pixar shouldn't be allowed to apply to Intel job postings, I am simply asking that you don't cold call our people." (Exhibit 4) Pixar raised particular concerns about a small team of employees – the "RenderMan" team. (*Id.*) RenderMan is Pixar's program for rendering 3D animation and visual effects, and the RenderMan team was working closely with Intel to optimize the performance of Pixar's software on Intel's architecture. The Pixar team had only 15 engineers with unique skill sets, and thus recruiting even a few people away from it could undermine the team's work. (*Id.*)

Pixar also expressed specific concerns that Intel was actively soliciting Pixar employees who had proprietary Pixar information, and that those employees might use Pixar's information in connection with a joint project between Intel and DreamWorks, an important Pixar rival. (Exhibit 5, Exhibit 6) Pixar's concerns about information transfer threatened future collaborative projects between Intel and Pixar. (Exhibit 7) One Intel manager described "run[ning] into a snag with Pixar" after several Pixar employees left for Intel and explained that Mr. Brandeau was "incandescent about the possibility of their IP being used for the benefit of his competition." (Exhibit 8) Intel managers working with Pixar engineers determined that "the only way we're going to get past this logjam is [to] agree that ... we will no longer target Pixar employees for recruitment." (*Id.*)

Keenly aware that Pixar could terminate or otherwise curtail the parties' collaborative efforts, Intel responded quickly, developing a unilateral policy not to engage in uninvited cold calling of Pixar employees the very next day, October 10, 2008. Mr. Gelsinger described to Mr. Brandeau at Pixar the unilateral policy that Intel had decided to implement in response to Pixar's complaint: "[W]e greatly value your partnership. . . . we've immediately directed our recruiters that Intel will not proactively pursue any Pixar employees going forward." (Exhibit 4) Mr. Brandeau replied to thank Mr. Gelsinger; he did not offer or promise anything to Intel in return.

As Exhibit 9 makes clear, the policy was not a hiring ban or even a general solicitation ban. It was not "a complete moratorium – people who approach us are fair game." Intel's Strategic Recruiting Manager for the Americas, Christina Dickenson, described the policy: "[W]e've been asked to avoid directly sourcing any Pixar employees. . . . It's causing problems for DEG [an Intel business group]. We're still ok to talk to them if they come to us, but don't cold call or target them directly." (Exhibit 10) As that description makes clear, Intel developed its policy to address a problem raised by a critical business partner that was threatening continued collaborative efforts between the parties. This was a *unilateral* business decision not to engage in recruiting by uninvited cold calling.

*B. Dell*

#### I. Scope of Joint Collaboration

Dell is one of Intel's leading customers. Intel and Dell work closely together on a wide variety of collaborative projects across their organizations. Indeed, in part because they are engaged in so many projects, Intel and Dell have partnered since 2007 on the Joint Innovation Center (JIC), a collaborative engineering program dedicated to developing Dell server platforms

that incorporate Intel components, as well as developing other products, software, and documentation. Both companies have invested tens of millions of dollars in the JIC and dedicated vast engineering, marketing, and sales resources to the collaboration. The JIC represents a formal organizing structure for the sort of projects the companies have continuously engaged in over the years. These ongoing efforts have resulted in several successful new Dell servers and many other jointly developed products.

Intel and Dell were involved in many close collaborations during the 2004-06 period, when the alleged agreement the Staff is investigating was in effect. For example, in 2005, Dell became the first OEM to incorporate Intel's dual core processors into its computers, and Dell worked "closely with Intel on the new technology for several years."<sup>7</sup> Intel and Dell worked jointly with Schlumberger from 2005 to develop a new platform specifically for the oil and gas industry.<sup>8</sup> Intel and Dell jointly developed new solutions for enterprise software in 2005.<sup>9</sup> And Intel and Dell collaborated in 2006 on technologies to improve performance and reduce power consumption in servers, resulting in a 30 percent performance improvement with a 40 percent reduction in power use.<sup>10</sup>

Moreover, Intel and Dell have continuously partnered on projects extending far beyond engineering. This is reflected in the document attached as Exhibit 11, in which Dell describes its relationship with Intel. Dell discusses "mutual opportunities to collaborate" including joint marketing efforts, a "joint plan to address the next billion PC users," improving customer experience, joint retail marketing, and other efforts. Dell also discusses "mutual operational

<sup>7</sup> Exhibit 12, <http://content.dell.com/us/en/corp/d/press-releases/2005-02-08-03-intel-dual-core.aspx>.

<sup>8</sup> Exhibit 13, <http://www.intel.com/pressroom/archivc/releases/2005/20051109corp.htm>.

<sup>9</sup> Exhibit 14, <http://content.dell.com/us/en/corp/d/press-releases/2005-09-21-01-oracle-database-10g.aspx>.

<sup>10</sup> Exhibit 15, <http://content.dell.com/us/en/corp/d/press-releases/2006-09-12-00-strategy.aspx>.

CONFIDENTIAL BUSINESS INFORMATION

opportunities” which include “align[ing] Intel technology with Dell innovative products” and “align[ing] supply chains and fulfillment models.” At least three times a year, for the launch of new Dell products with Intel components in the desktop, notebook, and server areas, Dell and Intel engineers needed to work together to ensure that Intel’s components – including microprocessors, chipsets, motherboard, and wireless chips – functioned properly in the Dell system, and that Intel’s components functioned to support the features that Dell wished to offer. For example, as described in Exhibit 16, Intel needed “Dell’s help” in identifying Type-3 bugs to validate chipset performance in Dell machines.

To facilitate this collaboration, Intel and Dell employees co-located at the offices of each company and often worked together in a firewalled room in a laboratory at Intel’s headquarters in Oregon. Around 200 Intel engineers are typically dedicated to working on Dell products. Because both the systems and the software tend to be quite complicated, the collaboration between Intel and Dell requires engineers to exchange detailed, proprietary information about the products on which they are working. After the two companies’ engineers, working collaboratively, ensured that the companies’ products would operate together seamlessly, joint marketing, manufacturing, and supply chain teams facilitated the most efficient product launch and sale. Dell and Intel teams conducted joint quarterly business reviews. *See, e.g.*, Exhibit 17, which identifies as a highlight for the quarter “winning deals together” and the building of “joint pipelines.”

## 2. Intel’s Alleged Agreement with Dell

In February 2005, Paul McKinnon, Dell’s then-Vice President of Human Resources, complained to Intel about Intel’s recruiting of Dell employees, stating that “at this point, we are,

as a company, rethinking our arrangement with you.” (Exhibit 18) We understand that Dell was concerned that Intel would specifically target for hiring Dell employees that it came to know through the parties’ joint collaborate efforts, many of whom could have Dell trade secrets that Dell feared could find their way to Dell’s competitors through collaborations between Intel and those competitors. Intel, in turn, worried that Dell’s concerns could impair the companies’ close collaborative working relationship.

To enable the collaboration to continue efficaciously, the two companies essentially agreed to restrict the use of one method of soliciting each other’s employees – neither would engage in uninvited cold calling.<sup>11</sup> They never agreed (i) not to solicit employees who had indicated in some way an interest in considering other employment opportunities (such as an inquiry from the employee, posting by the employee on a national job board, or by networking); (ii) not to interview each other’s employees; or (iii) not to hire each other’s employees. Intel also agreed that, before extending a formal employment offer, Intel staffing personnel would ask Dell candidates if they had informed their Dell manager that they intended to leave Dell and would document that they had done so. Dell’s documents show that the alleged agreement ended in January 2006. Intel’s documents confirm that the agreement ended in early 2006. (Exhibit 19, Exhibit 20)<sup>12</sup>

Notwithstanding this narrow and short-term alleged agreement, as Intel’s documents produced to the Division make clear, Intel regularly recruited and hired Dell employees. For example, Richard Taylor, an Intel Vice President for Human Resources, determined that Intel

<sup>11</sup> One document suggests that Michael Dell and Paul Otellini agreed not to engage in cold calling of employees from each other’s companies. (Exhibit 22) Our investigation has revealed no evidence that Mr. Dell and Mr. Otellini ever had any such discussions, and substantial evidence that no such discussions ever took place. We believe that the author of this e-mail was simply mistaken.

<sup>12</sup> Intel appears to have unilaterally continued its practice against uninvited cold calling of Dell employees. See, e.g., Exhibit 23.

CONFIDENTIAL BUSINESS INFORMATION

received over 1,700 unsolicited resumes from Dell candidates in 2005, and Intel hired 44 Dell candidates worldwide that year. (Exhibit 21) Intel hired applicants from Dell at more than twice its rate of hiring applicants from all companies. This belies any notion that Intel was boycotting or not aggressively seeking to hire Dell employees who expressed any interest in working for Intel.<sup>13</sup>

### C. Google

#### 1. Scope of Joint Collaboration

Intel's CEO, Paul Otellini, has served on Google's Board of Directors since April 2004. As with Pixar, Google and Intel have an extensive history of collaborating on joint development projects, which has been facilitated by Mr. Otellini's Google Board service. Intel established a "Google Program Office" in September 2008 to help manage the wide and continuing array of projects the companies are pursuing together. Renee James, a senior Intel executive, described the Google Program Office as a "process of improving our overall engagement and strategic dialog with this important ecosystem partner." (Exhibit 24)

Intel has long collaborated closely with Google to optimize the performance of Intel microprocessors that Google uses in its server farms, with Intel and Google employees working together in person and otherwise for long periods of time. Intel developed technology that helped Google translate speech in videos to text to enable Google's new video search product.<sup>14</sup> Intel is currently collaborating with Google on GoogleTV, Google's Chrome browser,<sup>15</sup> and the

<sup>13</sup> Intel data indicate that in 2007, the earliest year for which it has systematic data, its hiring rate for experienced external candidates in the United States was 1.07 percent. The hiring rate from unsolicited Dell resumes in 2005 was about 2.5 percent.

<sup>14</sup> See Exhibit 25, <http://www.intel.com/pressroom/archive/releases/2006/20060105corp.htm>.

<sup>15</sup> See Exhibit 26, [http://news.zdnet.com/2100-9595\\_22-343152.html](http://news.zdnet.com/2100-9595_22-343152.html).

Android operating system.<sup>16</sup> These projects require broad collaboration between engineering teams and other Intel and Google employees. For example, according to the Program Manager for GoogleTV, Google devotes fifteen engineers exclusively to the collaboration with Intel and calls on around 1,000 more engineers from the Android and Chrome projects to support the team. The GoogleTV team meets on a weekly basis. In the marketing area, Intel and Google evaluated a potential program to combine the strength of Intel's direct sales force with Google's plan to offer the Google toolbar for company-specific applications. Intel has worked with Google to optimize the performance of Google's software on Intel's architecture. Google relies heavily on Intel's open source team for various Chrome needs, such as graphic systems. Intel and Google have also collaborated on WiMax technology generally,<sup>17</sup> and specifically in developing WiMax offerings and applications for Nokia and Motorola.<sup>18</sup> The results of these close collaborations will be products that reach the marketplace starting in 2010 that will give consumers cutting edge capabilities, such as the merging of Internet and TV functionality.

## 2. Alleged Intel/Google Agreement

In 2007, Google began to initiate contact with Intel engineers to solicit them for positions at Google. Specifically, Google had solicited Dr. Brad Chen and David Sehr, members of Intel's compiler team. Google was successful in recruiting both engineers. This heightened Intel's concern that it might lose to Google those employees who were involved in joint development efforts. Accordingly, in September 2007, Mr. Otellini notified Google that he was concerned

<sup>16</sup> See Exhibit 27, <http://www.geek.com/articles/mobile/intel-seeking-googles-android-os-for-future-mids-20090710/>.

<sup>17</sup> See Exhibit 29, <http://www.techdirt.com/articles/20080506/1740131048.shtml>; Exhibit 30 <http://arstechnica.com/old/content/2008/03/google-intel-and-cable-companies-ponder-nationwide-wimax-jv.ars>.

<sup>18</sup> See Exhibit 31, [http://www.networkrevolutionvision.com/newt/l/wimaxvision/article\\_view.html?artid=20017465552](http://www.networkrevolutionvision.com/newt/l/wimaxvision/article_view.html?artid=20017465552); Exhibit 32.



Google was actively soliciting Intel employees. Eric Schmidt, Google's CEO, responded by saying that he understood Intel's concern and did not want to undermine the trust on which the Google/Intel collaborations are based, observing that "[w]e take these relationships exceptionally seriously." (Exhibit 28)

Mr. Otellini's emails indicated he believed he had an "unofficial no poaching policy" with Google and that he had a "handshake" agreement with Google not to poach (*i.e.*, to engage in uninvited cold calling) each other's senior talent.<sup>19</sup> (Exhibit 32, Exhibit 33) His comments as a layperson, however, should not be confused with legal conclusions about the nature of any communications he had with Google. Because Intel intended only to raise concerns with Google about Google's uninvited cold calling, Intel did not take any steps formally to implement any such "agreement" not to "poach" from Google. Intel did not train any recruiters or other employees not to cold call or otherwise actively recruit Google employees. Indeed, our investigation has revealed that Intel's Intel staffing employees generally had no knowledge about any such alleged agreement. Intel also understands that in 2005, Google had instituted a policy under which it would not engage in uninvited cold calling to solicit Intel employees because of the large number of joint collaborative projects between the companies, because Intel's CEO, Paul Otellini, served on Google's Board, and, relatedly, because Mr. Otellini's board service had allowed Mr. Otellini to identify a large number of opportunities for the two companies to collaborate on various projects. Intel understands that Google's policy was limited to uninvited

<sup>19</sup> Indeed, although for the purposes of this submission, Intel is focused on the factual and legal reasons why this alleged agreement with Google must be judged under the rule of reason, Intel does not believe there was any actual agreement between it and Google. Intel's produced documents indicate as much. For example, an email Mr. Otellini sent internally on September 6, 2007, mentions a "'no recruit'" between the two companies (Exhibit 32), but only three weeks later he and Mr. Schmidt exchange emails in which he seeks clarification as to what Google's policy is. (Exhibit 28) If the two CEOs had reached a meeting of the minds and formed an agreement, Mr. Otellini would not have needed to ask Mr. Schmidt to clarify Google's policy for him – and Mr. Schmidt would not have needed to "check" with someone internally at Google regarding Google's policy not to cold call Intel employees.

CONFIDENTIAL BUSINESS INFORMATION

cold calls; any Intel employee who expressed any interest in leaving Intel – including applying for a job at Google or mentioning his/her interest to a friend at Google – could be actively recruited and hired without violating Google’s internal policy. Mr. Otellini was aware that Intel employees were not barred from moving to Google. (Exhibit 33)

**II. The Alleged Agreements to Restrict Uninvited Cold Calls Are Ancillary to Broader Collaborations and Must Be Assessed Under the Rule of Reason.**

Intel’s alleged agreements with Pixar, Dell, and Google are ancillary to Intel’s broader business arrangements with those companies and, as such, must be assessed under the rule of reason. In addition, even if the Division were to take the position that the alleged agreements are not strictly ancillary, the rule of reason – not the per se rule – is plainly the proper legal framework for assessing them.

*A. The Rule of Reason Applies to Agreements Ancillary to Cooperative Ventures.*

Under the ancillary restraints doctrine, agreements that are ancillary to broader business ventures are evaluated under the rule of reason – even if they would be deemed illegal per se if not ancillary. “[A]ncillary restraints . . . are part of a larger endeavor whose success they promote.” *Polk Bros., Inc. v. Forest City Enterprises, Inc.*, 776 F.2d 185, 188-89 (7th Cir. 1985). “A restraint is ancillary when it may contribute to the success of a cooperative venture that promises greater productivity and output.” *Id.* at 189; *see also Valley Drug Co. v. Geneva Pharms., Inc.*, 344 F.3d 1294, 1313 n.31 (11th Cir. 2003) (“Agreements that are anticompetitive when considered in isolation. . . can still be lawful if they are ancillary to another agreement and, when viewed in combination, will have the overall effect of enhancing competition.”). *See also* 2000 Department of Justice and Federal Trade Commission Antitrust Guidelines for Collaborations Among Competitors § 3.2.

Courts have consistently held that ancillary restraints are judged under the rule of reason. *Sullivan v. National Football League*, 34 F.3d 1091, 1102 (1st Cir. 1994); *Los Angeles Memorial Coliseum Comm'n v. National Football League*, 726 F.2d 1381 (9th Cir. 1984); *Schering-Plough Corp. v. F.T.C.*, 402 F.3d 1056 (11th Cir. 2005); *Polk Bros.*, 776 F.2d at 189; *Andrx Pharms., Inc. v. Biovail Corp. Int'l*, 256 F.3d 799, 811 (D.C. Cir. 2001) (“ancillary restraints that facilitate productive activity are not . . . unlawful *per se*”).

The Staff has told us that they view the ancillary restraints doctrine as requiring a showing that the restraint was strictly, or at least reasonably, “necessary” to facilitating productive activity. Although the alleged agreements at issue here satisfy that standard, we believe the standard is overly narrow. While this standard may define the boundaries of safe harbors under the Competitor Collaboration Guidelines, it does not define the boundaries of the law, which merely requires that the restraint is “capable” of enhancing efficiency or makes the joint activity “more efficient.” *See, e.g., Rothery Storage & Van Co. v. Atlas Van Lines, Inc.*, 792 F.2d 210, 229 (D.C. Cir. 1986) (“an ancillary horizontal restraint, one that is part of an integration of the economic activities of the parties and appears capable of enhancing the group’s efficiency, is to be judged according to its purpose and effect”); *Tower Air, Inc. v. Fed. Express Corp.*, 956 F. Supp. 270, 283 (E.D.N.Y. 1996) (“Thus, to be lawful, the ancillary restraint must be subordinate to a separate, legitimate transaction and serve to make the main transaction more effective in accomplishing its purpose.”); XI Phillip E. Areeda & Herbert Hovenkamp, *Antitrust Law* ¶ 1912 at 320 (“An ancillary restraint is one that is reasonably related to a joint venture or transaction that, at least upon initial examination, promises to increase output, reduce costs, improve product quality, or otherwise benefit consumers. As a result, the profitability of an

ancillary restraint depends not on the market power of its participants but on their ability to reduce their costs or increase demand for their product.”).

*B. The Pixar, Dell, and Google Alleged Agreements Are Ancillary to Broader Business Collaborations.*

Intel’s alleged agreements with Pixar, Dell, and Google are “ancillary restraints that facilitate productive activity,” *Andrx Pharmaceuticals*, 256 F.3d at 811, and therefore must be evaluated under the rule of reason. Intel has deep, continuing collaborative relationships with Pixar, Dell, and Google. Intel works closely with these companies at various organizational levels and across departments on a continuing array of joint projects. As set forth in Part I, Intel entered into the alleged agreements or adopted unilateral policies against uninvited cold calling to promote these pro-competitive collaborations – not to restrict competition in labor markets – and the alleged agreements or policies delivered the desired collaborations.

*i. Pixar*

As discussed above, Intel unilaterally adopted a policy not to employ one method of solicitation – uninvited cold calling for Pixar employees. It did not reach any agreement with Pixar not to do so. Independent actions, of course, fall outside the purview of Section 1. *See Monsanto Co. v. Spray-Rite Serv. Corp.*, 465 U.S. 752, 761, 763 (1984) (warning that inferring a conspiracy from a complaint can “deter or penalize perfectly legitimate conduct”). The mere fact that Intel chose to respond to Pixar’s complaints by adopting a policy not to cold call Pixar employees did not create an agreement for Section 1 purposes. *See id.* at 752, 764 n.9 (policy decision in response to complaint does not constitute agreement absent evidence that party “communicated its acquiescence or agreement, and that this was sought by the [other party]”). Nor did Intel’s unprompted communication of its policy to Pixar create an agreement. *Id.*; *See*

*also Am. Council of Certified Podiatric Physicians & Surgeons v. Am. Bd. of Podiatric Surgery, Inc.*, 185 F.3d 606, 620 (6th Cir. 1999) (evidence showing that parties were “on friendly terms” and shared information about policies did not meet evidentiary burden to show evidence tending “to exclude the possibility of independent action”). Moreover, Pixar’s representative, Mr. Brandeau, did not promise anything in response to Mr. Gelsinger’s representation to him; he did not, for example, commit that Pixar would continue its collaborative efforts or take any other action suggesting Intel and Pixar “had a conscious commitment to a common scheme.” *Monsanto*, 465 U.S. at 768.

But, even assuming an agreement, this would represent a classic ancillary restraint. Intel and Pixar were working jointly on rendering technology, and in response to concerns from Pixar, Intel decided not to solicit Pixar employees. As discussed above, Pixar was especially concerned that Intel might solicit away Pixar employees who had proprietary information, particularly given Intel’s collaborative relationship with Pixar’s rival Dreamworks. That concern threatened to cause Pixar to withdraw from or restrict its support of the collaboration.

Indeed, Pixar explicitly stated that Intel’s recruiting of Pixar employees could jeopardize joint development projects. An Intel employee reported on a conversation with a Pixar complainant: “He is concerned about *future work with Intel on Larrabee visualization*. He wants to work with SSG [an Intel group] but [is] concerned [the] same thing [IP transfer from employees working on joint projects] might happen again.” (Exhibit 7, emphasis in original) *See also* Exhibit 34 (email from an Intel employee on the Pixar account team stating “thank you ... for helping resolve this issue in a quick and efficient manner! Much appreciated by the

account team so hopefully now we can focus on more productive opportunities with Pixar moving forward.”).

Therefore, if the parties could be deemed to have entered an agreement, the agreement furthered the parties’ collaborative efforts. Under the cases, it is thus ancillary to their collaboration.

ii. *Dell*

The importance of Intel’s relationship with Dell, its close collaborator and one of its largest customers for many years, cannot be overstated. Dell’s and Intel’s joint projects would have been jeopardized if concerns about Intel making uninvited cold calls to Dell employees it encountered through collaborations, employees who had not shown any interest in leaving their employment at Dell, caused Dell to avoid partnering or to be less open in its working relationships with Intel (or vice versa). Indeed, historical documents reflect concerns at both Intel and Dell that tensions over direct employee recruitment were threatening joint Intel-Dell collaborations. Intel’s Director of Enterprise Server Development, Eric Hooper, wrote an email to others at Intel stating that “[i]t is very important to our business and our relationship with Dell that we do not create undesired turnover.” (Exhibit 35) Then General Manager of the Server Platforms Group, Diane Bryant, responded that “[w]e definitely don’t want to impact the positive relationship you’ve built.” (*Id.*) In another email, Dell’s HR Vice President wrote “our management team is starting to escalate this,” referring to recruitment issues. (Exhibit 21) These are precisely the type of concerns that an agreement to refrain from uninvited cold calling among collaboration partners is designed to address.

Relatedly, as close business partners, both Intel and Dell have a legitimate interest in not undermining their partnership in the many key areas in which they collaborate by enticing critical employees, who have not otherwise expressed an interest in departing, to leave. As one court has put it, “[e]njoy[ing] the services of the experienced [] employees” is an “obviously sound business purpose” for an ancillary agreement. *Cesnik v. Chrysler Corp.*, 490 F. Supp. at 859, 866 (D. Tenn. 1980). And “retain[ing] the skilled services of management-level employees” to avoid “disruption of the ongoing business” is a legitimate, pro-competitive purpose. *Id.* at 868.

iii. *Google*

Intel and Google’s collaborative efforts involve very close relationships and sharing of substantial confidential information among key employees across a broad range of the two companies. These productive working relationships are enhanced when the companies are able to work closely together without fear that their key employees will be recruited away based on relationships that develop during their collaborative efforts. Like the situation with Pixar and Dell, active solicitation of key employees not looking for employment elsewhere could seriously undermine the success of joint collaborative efforts between Intel and Google and deprive consumers of the benefits that result from those efforts. In addition, Mr. Otellini’s service on Google’s Board of Directors brings pro-competitive benefits to both Google and Intel and, ultimately, to consumers. Google gains the board services of an experienced CEO from a large, highly respected, innovative technology company; Mr. Otellini brings expertise in developing strategy, managing a large organization, negotiating complex deals, and running sophisticated business operations, among other areas. Mr. Otellini’s Board service facilitates identifying a

wide range of projects on which the two companies can collaborate productively – to the benefit of both companies and consumers.

*C. The Alleged Agreements Were Not Overbroad.*

We understand that the Staff may have concerns whether the alleged agreements not to engage in uninvited cold calling with Pixar, Dell, and Google were overbroad because they were not limited by their terms to specific employees. Those concerns are misplaced. To qualify for ancillary treatment, a restraint need not be the least restrictive theoretical option that might have been available to the parties; rather, the courts look to whether the parties could *practicably* have obtained the benefits from the restraint through an agreement that is *substantially less restrictive*. Here, it would have been infeasible for the parties to implement a more narrowly targeted agreement that would have been substantially less restrictive.

To qualify for ancillary treatment, a restraint need not have been perfectly tailored to achieve its pro-competitive benefits; rather, the relevant inquiry is “whether the restraint. . . *goes so far beyond* what is necessary as to provide a basis for the inference that its real purpose is the fostering of monopoly.” *Syntex Labs., Inc. v. Norwich Pharmacal Co.*, 315 F. Supp. 45, 56 (S.D.N.Y. 1970) (emphasis supplied); *see also Rothery Storage & Van Co. v. Atlas Van Lines, Inc.*, 792 F.2d 210, 224 (D.C. Cir. 1986) (“the restraint imposed must be *related* to the efficiency”) (emphasis supplied); *Newberry v. Washington Post Co.*, 438 F. Supp. 470, 475 (D.D.C. 1977) (holding that defendant need not show that the approach chosen was the “least restrictive alternative” in a rule of reason case); *Transamerica Computer Co. v. IBM Corp.*, 481 F. Supp. 965, 1022 (N.D.Cal. 1979) (“[i]t is the choice of an unreasonable alternative, not the failure to choose the least restrictive alternative, that leads to liability”).



To evaluate whether an ancillary restraint is overbroad, the courts ask “whether comparable benefits could be achieved through a substantially less restrictive alternative.” *U.S. v. Brown Univ.*, 5 F.3d 658, 679 (3d Cir. 1993) (applying rule of reason in a case brought by the Division). *See also Am. Motor Inns, Inc. v. Holiday Inns, Inc.*, 521 F.2d 1230, 1248-49 (3d Cir. 1975) (“[I]n a rule of reason case, the test is not whether the defendant employed the least restrictive alternative.”). In this area, the Division’s Competitor Collaboration Guidelines are consistent with the case law. They state that “the Agencies do not search for a theoretically less restrictive alternative”; instead, they examine whether there are alternatives that are “practical” and “significantly” less restrictive.<sup>20</sup>

1. There Is No Practical Way for the Parties to Define a Narrower and Significantly Less Restrictive Agreement Not to Engage in Uninvited Cold Calling.

There were no alternatives to Intel’s alleged agreements with Dell and Google that are both “practical” and “significantly” less restrictive, as discussed below. As for Pixar, it has not been a significant source of employees for Intel at any time; Intel’s records indicate that it hired fewer than three employees from Pixar in any of the years from 2005 to 2009. The relatively few Pixar employees that Intel might have been interested in soliciting were the sort of technical employees who were involved in the joint collaborations; Intel would generally have little interest in soliciting most other employees from Pixar (*e.g.*, animators or writers). Accordingly, as a practical matter, the alleged agreement with Pixar was not overbroad because it affected only employees engaged in collaborative efforts with Intel.

*a. Dell*

<sup>20</sup> 2000 Department of Justice and Federal Trade Commission Antitrust Guidelines for Collaborations Among Competitors § 3.36(b).

CONFIDENTIAL BUSINESS INFORMATION

Intel's alleged agreement with Dell is not specifically limited to certain types of employees. The joint projects between Intel and Dell were ever-changing as projects were added, terminated, and revised; and they involved a continually changing roster involving thousands of employees. There is no practical way that the parties could have defined in advance which employees should be covered by an agreement not to engage in uninvited cold calling.

Notably, Intel responded to Dell's complaints by taking those measures that were necessary to address Dell's concerns and allow the parties' many collaborations to proceed effectively. Those measures, of course, did not address every conceivable issue that might come up with respect to recruiting or proprietary information transfer between the Intel and Dell. Rather, the two companies narrowly tailored any agreement between them to address particularly acute concerns regarding uninvited cold calling without unduly interfering with the mechanisms by which their employees transfer between the companies.

The alleged Dell agreement included an additional feature: Intel hiring personnel were to ask Dell employees whether they had informed their managers that they were seeking employment with Intel. This requirement served two purposes. First, it reduced the likelihood that Dell employees would leave without giving any advance notice – a problem that had frustrated Dell and created tensions in the working relationship between Dell and Intel. *See, e.g.*, Exhibit 36. Second, it gave Dell a chance to negotiate with the employee, and, if desired, perhaps match the offer expected from Intel or to find other ways of encouraging the employee to stay with Dell (*e.g.*, reassignment to a new position or a new title). Both of these purposes

served to avoid disputes and further collaborative efforts between the parties, and the latter generally benefitted the allegedly affected employees.

*b. Google*

We do not believe that Intel and Google reached any agreement regarding cold calling of employees. But even if one assumed an agreement, that it was not specifically limited to certain types of employees does not suggest that it was overbroad.

It would be highly impractical for Intel and Google to determine in advance which employees (i) might be participating in the multiple ongoing cooperative projects involving both companies or (ii) might come to Google's or Mr. Otellini's attention through his Google Board service. Given the breadth of the cooperative activities between Intel and Google, it would be difficult or impossible to define in advance a limited group of employees to identify in an "agreement" not to engage in uninvited cold calling. Similarly, Google's Board regularly addresses topics ranging from marketing to finance to strategy to operations and more. Any of these topics might involve presentations by Google employees or comments by Mr. Otellini that draw on his experiences at Intel. That being so, there is no practical means by which Intel and Google could specify in advance the particular employees that should be covered by an agreement not to engage in uninvited cold calling.

2. Narrower Agreements Not to Engage in Uninvited Cold Calling Would Not Have Been Administrable and Would Have Led to Increased Litigation Risk Between Collaborators.

Intel currently employs over 80,000 people throughout the world. Its staffing organization received resumes from over 177,000 individuals in 2009, and over 275 staffing employees in many different locations managed the recruitment and hiring process. (During

periods of intensive hiring, Intel has had a much larger staffing organization.) These staffing employees range from the highly sophisticated to those with lower skill levels, and there is a high level of turnover in some parts of Intel's staffing organization. Any agreements that Intel enters into must be effectively administrable by this diverse group of employees.

Any agreement or unilateral policy not to engage in uninvited cold calling with Dell or Google limited to a narrower range of employees would be impractical to administer. For example, Intel's staffing organization would be required to keep tabs on all of the collaborative projects with Intel and Dell, some of which are confidential. Intel would have had to develop a mechanism for Intel employees working with Dell continuously to report to Intel's worldwide staffing operations about which Dell employees were involved with an Intel collaboration at the moment – an extremely difficult task given the breadth of Intel's and Dell's collaborative activities and the uncertainty inherent in defining which Dell employees would need to go on a no uninvited cold calling list. (There would be similar barriers to defining more narrowly a group of Intel employees that Dell would be prohibited from actively soliciting.) Any attempt to maintain an employee-by-employee no cold calling list for Google would run into similar impediments, creating an administrative difficulty that is far disproportionate to the minimal role cold calling plays in Intel candidate sourcing. The law does not require companies to implement alternatives that are overly complicated and impractical to administer. *See, e.g., Buffalo Broad. Co., Inc. v. Am. Soc. of Composers, Authors & Publishers*, 546 F. Supp. 274, 286-87 (S.D.N.Y. 1982) (overruled on other grounds) (“a realistically available . . . alternative is one that is reasonably efficient and not unreasonably costly”).

Moreover, an agreement that cannot be accurately and effectively administered would lead to breaches, thereby creating tensions between the parties that undermine the very purpose of the alleged agreement. It is clear from the documents Intel has produced that it regularly faced complaints from Dell about recruiting – even with a relatively simple and limited non-solicitation agreement restricting uninvited cold calling in place. Had Intel implemented a more intricate agreement that carefully delineated precisely which Dell employees were not to be the targets of uninvited cold calls, the risk of error by Intel’s staffing organization would have been intolerably high. Such errors would have undermined the ability of Intel to collaborate with one of its most significant collaboration partners.

Finally, this is not a case of companies implementing overbroad restraints as a sham in order to conceal an actual intent to thwart competition. Unlike with many ancillary agreements – e.g., those involving agreements not to compete between competing suppliers in concentrated product markets – any over-breadth had no potential to harm competition in a properly defined market. As discussed below, these alleged bilateral agreements had no potential to injure competition in any relevant employment market, which consists of the dozens (if not hundreds) of employers. That being so, it is implausible that the real purpose – let alone effect – of the alleged agreements was “the fostering of monopoly,” *Syntex*, 315 F. Supp. at 56, thereby taking them out of the rule of reason.

### **III. The Rule of Reason Would Apply Even If the Alleged Agreements Were Not Deemed Ancillary Restraints.**

We understand that Staff is considering recommending that the Division challenge as per se illegal any agreement that it concludes does not fall within a narrow definition of an ancillary restraint. That approach would be contrary to the law.

To begin, the Staff's theory appears grounded in a misapprehension that there is a strict dichotomy between ancillary agreements and per se unlawful agreements. That is, the Staff seems to contend that if any agreement does not qualify for ancillary treatment, it must be illegal per se. But, as a leading commentator has put it, "[w]hile the most frequent application of the rule of reason involves restraints that are 'ancillary' to some underlying productive joint venture, a significant number of challenged agreements qualify for rule of reason treatment even though they are not realistically ancillary to anything." XI Phillip E. Areeda & Herbert Hovenkamp, *Antitrust Law* ¶ 1912(c)(3), at 289; *see also Major League Baseball Props., Inc. v. Salvino, Inc.*, 542 F.3d 290, 339 n.7 (2d Cir. 2008) ("a non-ancillary restraint is not necessarily unlawful or evaluated under a *per se* rule; rather it is simply evaluated independent of the joint venture. . ."); *U.S. v. Brown Univ.*, 5 F.3d 658 (3d Cir. 1993) (applying rule of reason to financial aid agreement); *Vogel v. Am. Soc'y of Appraisers*, 744 F.2d 598, 603 (7th Cir. 1984) (applying rule of reason given "novelty of the challenged practice").

The Supreme Court has made clear that the rule of reason is the "prevailing standard" for assessing the legality of a given agreement, *Sylvania*, 433 U.S. at 36, 49 (1977), and that it "presumptively applies rule of reason analysis" to Section 1 cases, *Dagher*, 547 U.S. at 5. *See also Leegin Creative Leather Prods., Inc. v. PSKS, Inc.*, 551 U.S. 887, 885 (2007) ("The rule of reason is the accepted standard for testing whether a practice restrains trade in violation of § 1").

It is the per se rule, not the rule of reason, that is narrowly confined to specific circumstances. As the Second Circuit put it recently, there is only a "narrow range of behavior that is considered so plainly anti-competitive and so lacking in redeeming pro-competitive value" that it should be per se illegal. *Geneva Pharms. Tech. Corp. v. Barr Labs. Inc.*, 386 F.3d 485,

506 (2d Cir. 2004). To justify a per se prohibition, a restraint must “lack . . . any redeeming virtue,” *Northwest Wholesale Stationers v. Pacific Stationery & Printing Co.*, 472 U.S. 284, 289 (1985), and have “manifestly anti-competitive” effects, *Continental TV Inc. v. GTE Sylvania, Inc.*, 433 U.S. 36, 50 (1977). *See also Illinois ToolWorks Inc. v. Independent Ink, Inc.*, 547 U.S. 28, 29 (2006). For the reasons discussed below, the alleged agreements not to solicit by uninvited cold calling at issue here are plainly not “manifestly anticompetitive.” And because they were in furtherance of pro-competitive business collaborations, they have “redeeming virtues.”

Moreover, the “*per se* rule is appropriate only after courts have had considerable experience with the type of restraint at issue . . . and only if [court]s can predict with confidence that [it] would be invalidated in all or almost all instances under the rule of reason . . .” *Leegin*, 551 U.S. at 886-87. Courts do not have “considerable experience” with agreements not to engage in uninvited cold calling. Indeed, courts reviewing markedly more restrictive no-hire agreements have concluded time and again that the effects of no-hire agreements on competition are far from clear, cannot be presumed, and thus must be tested under the rule of reason.

For example, in *Union Circulation Co. v. F.T.C.*, 241 F.2d 652 (2d Cir. 1957), the Second Circuit applied the rule of reason to a *multilateral no-hire* agreement among several magazine subscription sales agencies. This agreement was a flat hiring ban (*i.e.*, broader than a non-solicitation agreement), was not ancillary to any business arrangement among the agencies, and involved “leaders in the field [that] constitute a very substantial segment of the industry.” *See id.* at 654, 656. The court applied the rule of reason and specifically rejected per se treatment: “[b]ecause a harmful effect upon competition is not clearly apparent from the terms of

these agreements, we believe them to be distinguishable from those boycotts that have been held illegal *per se*.” *Id.* (citation omitted).<sup>21</sup>

The Seventh Circuit has also spoken to the issue in the context of an agreement among encyclopedia sales organizations. The Court applied a “standard of reasonableness” despite the restraint entirely preventing the hiring of former employees for six months after termination. *Nichols v. Spencer Int’l Press, Inc.*, 371 F.2d 332, 337 (7th Cir. 1967). The court explained that the Supreme Court had “seldom considered, in light of the antitrust laws the validity of employer agreements directed solely at industry employment practices.” *Id.* It held that “[a]greements not to compete are tested by a standard of reasonableness.” *Id.*

The federal enforcement agencies typically take a similar approach. For example, in *U.S. v. Arizona Hosp. & Healthcare Assoc. & AzHHA Service Corp.*, the Department entered a consent decree with AzHHA regarding conduct including operating a joint registry for purchases of temporary nursing services that included maximum rates and other restrictions. *See Compl.*, No. CV07-1030 (D. Ariz. Sept. 12, 2007). The Department did not allege this behavior was *per se* unlawful, even though it included rate setting and far exceeded the conduct at issue here in terms of limiting employee mobility. *Id.*<sup>22</sup>

Courts and federal agencies acknowledge that the competitive impact of no-hire agreements is unclear and requires full consideration of the facts. That is all the more the case for non-solicitation agreements, which are far less restrictive than the no-hire agreements

<sup>21</sup> Forty years later, the Second Circuit again confirmed that the rule of reason applies to no-hiring agreements, this time in considering an agreement regarding insurance agents in *Bogan v. Hodgkins*, 166 F.3d 509, 515 (2d Cir. 1999): “Although such agreements fall within the ambit of antitrust law, we have nonetheless denied them *per se* treatment because ‘a harmful effect upon competition [was] not clearly apparent.’”

<sup>22</sup> The Department alleged *per se* illegality in only one case involving a non-solicitation agreement, *U.S. v. Assoc. of Family Practice Residency Dirs.* But that case also involved other conduct going far beyond a non-solicitation agreement; the Division also alleged overt price-fixing and market allocation, neither of which is present here. *See Compl.*, No. 96-575-CV-W-2 (W.D. Mo. May 28, 1996).

CONFIDENTIAL BUSINESS INFORMATION



examined in most cases. As one court explained, the key distinction between a no-hire and non-solicitation agreement is that a no-hire “on its face . . . forb[ids] the parties not only from active solicitation, but also from hiring any employee who left . . . on his or her own.” *Celtic Maint. Servs., Inc. v. Garrett Aviation Servs., LLC*, No. CV 106-77, 2007 WL 4557775, at \*3 (S.D. Ga. Dec. 21, 2007).

In short, uninvited cold calling agreements are not manifestly anticompetitive; have redeeming virtues, such as those described above regarding Pixar, Dell, and Google; and do not have anticompetitive effects. They are therefore subject to the rule of reason and cannot properly be condemned under the per se rule.

#### **IV. The Alleged Agreements Easily Withstand Scrutiny Under the Rule of Reason.**

The alleged agreements to restrict uninvited cold calling at issue here are plainly lawful under the rule of reason. As described below, the Division would need to show anticompetitive effects to prove a Sherman Act violation, and such a showing is implausible in a properly defined market. Additionally, available data confirm that the alleged agreements resulted in no anticompetitive effects. Finally, Intel’s recruiting efforts only infrequently rely on uninvited cold calling because it is both ineffective and fraught with legal risk, and bilateral restrictions therefore do not meaningfully affect even hiring by Intel or its counterparties to the alleged agreements, let alone injure competition in a properly defined employment market.

##### *A. The Division Bears the Burden of Proving Anticompetitive Effects.*

Under the rule of reason, the Division would bear the burden of proving anticompetitive effects in a properly defined employment market. *See, e.g., Clorox Co. v. Sterling Winthrop*, 117 F.3d 50, 59-60 (2d Cir. 1997) (absent showing of anticompetitive effect from relevant conduct,

“immaterial” whether conduct had pro-competitive benefits); *Schachar v. Am. Acad. of Ophthalmology, Inc.*, 870 F.2d 397, 400 (7th Cir. 1989) (once it is determined that defendant’s conduct brought no anticompetitive effects, that is the end of the rule of reason analysis); *Lektro-Vend Corp. v. Vendo Co.*, 660 F.2d 255, 269 (7th Cir. 1981) (“requir[ing] that the anticompetitive effect of an ancillary restraint be proven”); *U.S. v. Visa U.S.A., Inc.*, 163 F. Supp. 2d 322, 345 (S.D.N.Y. 2001) (“Under the rule of reason, the Government bears the initial burden (by a preponderance of the evidence) of demonstrating that each restraint has substantial adverse effects on competition.”).

Even under a “quick look” approach, a plaintiff must show that, “based upon economic learning and the experience of the market, it is obvious that a restraint of trade likely impairs competition.” *Polygram Holding, Inc. v. F.T.C.*, 416 F.3d 29, 36 (D.C. Cir. 2005). Here, there is no economic learning or experience of the market relevant to alleged bilateral agreements restricting one method of recruiting employees. *Cf. U.S. v. Microsoft*, 253 F.3d 34, 84 (D.C. Cir. 2001) (declining to condemn per se tying arrangements involving platform software products because there was “no close parallel in prior antitrust cases” and “simplistic application of per se tying rules carries a serious risk of harm”). And, as we have explained, there is in fact every reason to believe that the alleged “agreements” not to engage in uninvited cold calling had pro-competitive rationales and no potential to harm competition.

*B. The Alleged Agreements Have No Anticompetitive Effects in a Properly Defined Antitrust Market.*

The alleged agreements here have no effect on competition in a properly defined labor market. They do not significantly restrict the movement of employees – even between the parties – because they do not (i) prevent an employee from leaving his or her current employer

CONFIDENTIAL BUSINESS INFORMATION

for the other firm or (ii) prevent one from soliciting employees of the other who have in any way indicated an interest in seeking employment opportunities.

Moreover, it is simply implausible that such an alleged bilateral agreement among two companies (like those at issue here) would have anticompetitive effects in a properly defined employment market. The relevant markets in this context are far broader than just “employment at Intel and Pixar,” or “employment at Intel and Dell,” or “employment at Intel and Google.” Intel, Pixar, Dell, and Google, along with many other companies – technology and non-technology companies alike – all recruit engineers, marketers, financial experts, administrative staff, and other types of employees. Indeed, the courts have recognized the breadth of employment markets and rejected attempts to limit antitrust markets to the hiring of employees in a particular industry – let alone *hiring of employees between two companies* within an industry. *See, e.g., Eichorn*, 248 F.3d at 147-148 (rejecting plaintiff’s narrow product market and defining the proper market as “all those technology companies and network services providers who actively compete for employees with the skills and training possessed by plaintiffs”); *In re Compensation of Managerial, Prof’l, and Tech. Employees Antitrust Litig.*, No. 02-CV-2924, 2008 WL 3887619, at \*9 (D.N.J. Aug. 20, 2008) (granting defendants’ motion for summary judgment because, with respect to all of the named plaintiffs, “the [relevant] markets extend in different ways and to different degrees beyond the oil and petrochemical industry”).

Bilateral agreements between Intel and another company could not harm competition in a properly defined employment market. The parties to any such agreement – which would be just two of at least dozens of employers competing to hire employees – plainly lack monopsony power in such a market. *See Capital Imaging Assocs. v. Mohawk Valley Med. Assocs.*, 996 F.2d

CONFIDENTIAL BUSINESS INFORMATION

537, 547 (2d Cir. 1993) (“[I]t is clear that with such a *de minimis* market share, appellees lack the power to injure competition.”).<sup>23</sup> Indeed, we do not understand the Staff even to contend that Intel’s alleged bilateral non-solicitation agreements with Pixar, Dell, or Google actually injured competition in a relevant market.

We understand that the Staff may seek to analogize the alleged agreements at issue here to a price fixing agreement among horizontal competitors. But horizontal price fixing agreements do not arise in the context of legitimate, welfare-enhancing collaborations among the parties, and they almost always have the purpose and effect of reducing output or raising prices in the relevant market in which the parties compete. *See F.T.C. v. Superior Court Trial Lawyers Assoc.*, 493 U.S. 411, 432-33 (1990); *Bus. Elec. Corp. v. Sharp. Elec. Corp.*, 485 U.S. 717, 723 (1988) (“[P]er se rules are appropriate only for ‘conduct that is manifestly anticompetitive’”).

*C. Data Confirm that the Alleged Agreements Did Not Injure Competition.*

Intel has only limited data available about the number of employees that come from particular companies.<sup>24</sup> But that data suggest that these alleged agreements had no meaningful effect on hiring from those companies with which Intel allegedly had agreements. For example, Intel’s alleged agreement with Dell ended in January 2006. A contemporaneous email indicates that in 2005, when the alleged agreement was in effect, Intel [REDACTED]. Intel’s data from 2007, after the alleged agreement ended, indicates that Intel [REDACTED] [REDACTED] (comparable data are not available currently on a worldwide basis).

<sup>23</sup> See also *In re Compensation of Managerial, Prof'l & Tech. Employees Antitrust Litig.*, No. 02-cv-2924, 2008 WL 3887619 (D.N.J. Aug. 20, 2008) (granting summary judgment because plaintiff failed to define product market adequately for employees in wage information exchange litigation); *Haines v. Verimed Healthcare Network LLC*, No. 4:08cv791 CDP, 2009 WL 799583, at \*6 (E.D. Mar. 24, 2009) (employees have many employment options other than the two particular companies at issue in a non-compete).

<sup>24</sup> Intel began collecting these data in late 2006. The database frequently omits information regarding prior employers. But the lack of completeness would not tend to bias the data in any particular direction.

CONFIDENTIAL BUSINESS INFORMATION

Although Intel's overall hiring rate decreased between 2005 and 2007, these data are inconsistent with any notion that Intel's agreement with Dell not to engage in uninvited cold-calling had any material effect even on Intel's hiring of Dell employees – let alone on competition in a properly defined employment market.<sup>25</sup>

Dell maintains data showing companies from which it hired employees over a longer period. Dell's data show that it hires very few Intel employees and that the number of employees hired from Intel was small before, during, and after its agreement restricting uninvited cold calling of Intel employees:

2003
2004
2005
2006
2007
2008
2009

The Staff might suggest that agreements not to engage in uninvited cold calling could disadvantage an individual employee who does not receive a cold call as a result of such an agreement. But this ignores all the other methods by which an employee could find out about or take advantage of whatever employment opportunities might be available to him or her at the companies that are parties to the alleged agreement or elsewhere. These include contacting potential employers directly; indicating his or her openness to employment offers on a

25

This appears to be a function of Google employees' general reluctance to leave Google – given, among other things, Google's rapid growth and the opportunities that presents; and its reputation as an employer of choice. Although Intel does not track the companies to which its departing employees move, we understand that many Intel employees have moved to Google in the last few years.

CONFIDENTIAL BUSINESS INFORMATION

networking website, such as LinkedIn.com or Monster.com; retaining a headhunter; and receiving cold calls from other potential employers. In addition, even assuming an agreement not to engage in uninvited cold calling could harm a hypothetical individual who made no effort to take advantage of these alternatives, it would still not violate the antitrust laws because *competition* in a relevant market as a whole has not been injured. *See Clorox Co. v. Sterling Winthrop, Inc.*, 117 F.3d 50, 57 (2d Cir. 1997) (“to fulfill the requirement of showing an actual adverse effect in a [properly-defined] relevant market, ‘the plaintiff must show more than just that he was harmed by the defendant’s conduct’”) (internal citation omitted).

Finally, even if the Division could show anticompetitive effects, those effects would need to be balanced against the procompetitive benefits of the agreements under the rule of reason. As discussed in Part I, Intel’s alleged agreements with Pixar, Dell, and Google brought significant precompetitive benefits. Those benefits outweigh any conceivable anticompetitive effects that the Division might assert.

*D. Even Absent the Alleged Agreements, Intel Sparingly Used Uninvited Cold Calling as a Recruiting Practice; Accordingly, the Alleged “Agreements” to Limit Uninvited Cold Calling Could Not Have Had any Anticompetitive Effects.*

Any bilateral agreement not to engage in uninvited cold calling would have little if any effect on Intel’s hiring of employees from the counterparty to the alleged agreements because such cold calling plays only a very limited role in Intel’s recruiting efforts. Indeed, Intel’s 2009 data indicates that uninvited staffing department cold calling accounted for ***less than 2 percent of all Intel hires***. Of Intel’s total hires of 6,014 hires in 2009, 110 were originally contacted by Intel’s staffing department through cold calling. Intel sourcing specialists limit their use of such cold calling because of its ineffectiveness and the substantial risk of legal disputes. Instead, Intel

recruiters employ other readily available, more productive, efficient, and, as described below, less risky recruiting methods that focus on identifying qualified candidates who have already indicated an interest in considering new employment.

To meet its hiring demands, Intel generally seeks applicants who had previously indicated an interest in leaving their other employment. Since 2005, hundreds of thousands of applicants applied for positions at Intel. Intel uses a variety of recruiting methods to identify other interested potential candidates. Intel sourcing specialists constantly search national online job boards, such as Monster.com; CareerBuilder.com; LinkedIn.com; Jobserve.us; Dice.com; Witi.com (Women In Technology International); and Job.com. These boards post thousands of resumes that employers can search for qualified candidates. Additional recruiting methods include the Intel.com job site; referral incentives for current employees to recruit colleagues; attending job fairs; networking with other recruiters; and soliciting Intel employees for referral candidates. Employees also extensively network with colleagues at other employers and at conferences when interested in moving to another employer.

Given all of these other robust recruiting methods, Intel uses uninvited cold calling as a recruiting method only sparingly. Intel disfavors uninvited cold calling not only because it is inefficient and because sourcing specialists cannot quickly identify qualified applicants, but also because it increases the risk of intellectual property disputes.<sup>26</sup> Technology companies compete fiercely for qualified candidates, with open positions greatly exceeding the number of qualified

<sup>26</sup> Intel could also face employee raiding claims, which are those which arise out of the hiring practices or recruiting methods used by a potential employer.

CONFIDENTIAL BUSINESS INFORMATION

candidates available to fill them.<sup>27</sup> Employees frequently move between technology companies and projects taking the knowledge, skills and experience acquired at their previous jobs with them. In many instances, this knowledge includes a former employer's confidential intellectual property and know-how. This information is critical to the former employer's business success and was developed by that employer at enormous expense. Given that employees possess critical proprietary information, the magnitude of the investment made in developing confidential intellectual property, and the crucial role that confidential intellectual property plays in a technology company's success, employers often bring or threaten lawsuits to protect confidential intellectual property and to attempt to prevent their former employees from taking jobs with rival companies.<sup>28</sup> For example, Intel worked not only with Pixar but also with Dreamworks; as a result, Pixar's executives had substantial concerns that an Intel hire from Pixar would direct his or her insider knowledge towards projects with Dreamworks, a key Pixar competitor.

<sup>27</sup> As of November 17, 2009, Dice.com, which primarily advertises technical jobs, such as engineers, developers, and systems analysts, had more than 50,000 available jobs posted. On that same date, CareerBuilder.com offered more than 13,000 engineer positions, Monster.com showed more than 5,000 available engineer positions, and Jobserve.us displayed over 3,600 available engineer positions.

<sup>28</sup> In California, a center for technology companies, state law prohibits enforcement of non-competition agreements. Accordingly, employees often leave a company with critical property information and go to work for a rival.

CONFIDENTIAL BUSINESS INFORMATION



Virtually every company being investigated by the Division has been involved in high-profile disputes involving employee raiding claims.<sup>29</sup> The substantial volume of hiring at Intel of technical employees who typically have proprietary knowledge of their former employers makes the litigation risk a constant concern for Intel. Intel understandably wishes to avoid defending against costly, disruptive, and distracting employee raiding claims.

The most common employee raiding claims that former employers bring against firms that hire their former employees include unfair competition, tortious interference with contract, tortious interference with business expectancies, breach of contract (such as non-solicitation provisions in employment agreements), and related misappropriation of trade secrets. Targeting potential employees through uninvited cold calling increases a company's vulnerability to each of these causes of action. For example, uninvited cold calling leaves a hiring employer open to allegations that it identified the targeted employee because he or she holds key trade secrets belonging to the other company, which could be the basis of unfair competition and tortious interference with contract claims. By contrast, approaching a potential candidate who has already put him or herself on the job market blunts any allegations that an employee was targeted by the use of improperly obtained confidential information. Cold calling also increases an employer's exposure to misappropriation of trade secret claims because it raises suspicions and

<sup>29</sup> See *Johnson & Johnson Vision Care, Inc., v. Genentech, Inc.*, No. 06-cv-2898 (N.D. Cal. April 28, 2006) (Johnson & Johnson alleged that Genentech targeted and hired away seven key Johnson & Johnson employees in an effort to gain an "unfair competitive advantage and disrupt Johnson & Johnson's business."); *Motorola, Inc. v. Research in Motion Ltd.*, No. 08CH32748 (Ill. Cir. Ct., Cook Co. Sept. 5, 2008) (Motorola sued Research in Motion ("RIM") for violating a non-solicitation agreement between the parties in connection with RIM's targeted email campaign soliciting Motorola's employees); and *Microsoft Corp. v. Lee*, No. 05-2-23561 (Wash. Sup. Ct., King. Cty. Sept. 15, 2005) (enjoining Dr. Kai-Fu Lee, a former Microsoft employee, from recruiting Microsoft employees after Microsoft sued Google and Dr. Lee for breach of Dr. Lee's noncompete and nonsolicitation agreements); see also *Joost US, Inc. v. Volpi*, No. 1:09-cv-00708 (D. Del. Oct. 14, 2009) (Joost accused its former Chairman, Michaelangelo Volpi, of misappropriating confidential information, unlawfully soliciting away Joost employees, and breaching his fiduciary duty stemming from allegations that Mr. Volpi used his role as Chairman to solicit key Joost employees, about whom he had confidential information).

CONFIDENTIAL BUSINESS INFORMATION

permits allegations that an employee or group of employees has been targeted for the purpose of obtaining confidential intellectual property. Recruiting candidates who have previously expressed interest to the hiring employer deflates suspicions of an improper purpose and reduces the likelihood of allegations that the employer targeted the employee in order to misappropriate trade secrets. Accordingly, avoiding uninvited cold calling decreases the risk of having to defend employee raiding claims and related misappropriation of trade secret claims.

Indeed, Intel's own experiences demonstrate the risks of cold calling and why Intel minimizes the role of cold calling in its recruiting efforts. For example, in 1999, Motorola brought employee raiding claims against Intel after it used cold calling to recruit a group of microprocessor designers from a Motorola microprocessing design center. Motorola alleged that, in doing so, Intel must have used improperly obtained confidential personnel information from a former Motorola engineer who Intel had hired earlier. Although Intel had not done that, using cold calling as a recruitment method left Intel open to Motorola's accusations. In view of the litigation's significant potential distractions and disruption, Intel was forced to settle Motorola's claims.<sup>30</sup>

In 2006 and 2007, two Intel competitors separately threatened to sue Intel for improperly soliciting and hiring their employees, claiming tortious interference with contract, unfair competition, and misappropriation of trade secrets.<sup>31</sup> The competitors dropped their respective

<sup>30</sup> In 2006, Intel threatened to sue AMD for recruiting by cold calling and hiring more than ten Intel engineers, alleging that AMD would not have been able to target Intel's engineers had it not improperly obtained confidential personnel information. Intel and AMD settled this matter.

<sup>31</sup> In September 2006, ATI Technologies threatened litigation alleging that Intel improperly obtained and used confidential information, potentially including information that Intel obtained through collaborative activities with ATI to target ATI employees for recruitment with unsolicited emails and voice-mails. A copy of the threat letter from ATI is attached as Exhibit 37. In July 2007, Freescale Semiconductor accused Intel of misappropriating trade secrets and predatory hiring when Intel hired six engineers from Freescale. In both instances, the companies threatened claims of tortious interference with contractual relations, unfair competition, and misappropriation of trade secrets.

CONFIDENTIAL BUSINESS INFORMATION

claims when Intel responded that it had not in fact engaged in cold calling. In one case, Intel had relied on national posting sites to identify the recruited employees and, in the other, the recruited employee had contacted Intel himself.

The divergent outcomes of these matters shows the risks that cold calling creates for a technology company seeking to avoid employee raiding and related misappropriation of trade secret claims. Recruiting by uninvited cold calling leaves a technology company vulnerable to claims that it “targeted” potential employees through an improper method for an improper purpose. On the other hand, using recruiting methods directed at potential employees who have already indicated interest in considering other employment reduces the risk of these types of claims.

#### **V. Conclusion**

For the reasons set forth above, Intel’s alleged non-solicitation agreements with Pixar, Dell, and Google are ancillary to broader business ventures and must be evaluated under the rule of reason. Even if the alleged agreements were not deemed ancillary restraints, the rule of reason would still apply – the alleged non-solicitation agreements are plainly not the type of a restraint that, based on substantial judicial experience, always or almost always injure competition. Given that the agreements have no potential to harm competition in a relevant market and bring substantial pro-competitive benefits, they easily pass muster under the rule of reason. The Division should close its investigation without taking action.

# EXHIBIT 1



[Home](#) > [Press Room](#) > [News Releases](#) >

#### news resources

[Press Releases](#)

[Press Kits](#)

[Photos](#)

[Chip Shots](#)

[Innovation@Intel](#)

[Video / Audio](#)

[Media Center](#)

[Intel Blogs](#)

[Contact Intel PR](#)

#### Corporate Information

[About Intel](#)

[Biographies](#)

#### Legal Information

[Competition in the Innovation Economy](#)

#### Search News Resources

[Advanced Search >](#)

## Intel News Release

### Intel Helps Disney-Pixar Cook up a Blockbuster With 'Ratatouille'

#### From Movie-Making to Movie-Watching, Intel Multi-Core Processor Technology is a Key Ingredient

#### Related links

[More in this category](#)

[Corporate Information](#)

[Contact Corporate Press Relations](#)

SANTA CLARA, Calif., July 3, 2007 – Intel Corporation joined forces with Walt Disney Pictures and Pixar Animation Studios on the animated-comedy, "Ratatouille," which debuted in theaters on June 29 as the No. 1 movie in America. Intel-based computers have been a key technology ingredient in the digital media industry for more than 15 years, and for the first time the company is working hand-in-hand with Disney-Pixar® as a major promotional sponsor.

"Ratatouille" follows the adventures of a rat named Remy, who dreams of becoming a master chef in the culinary capital of Paris, boasts key breakthroughs in animation. The technology used to create and render the movie, Intel® Xeon™ processors with Intel® Core™ microarchitecture, helped build the emotional and visual excitement by delivering an experience packed with advances in computer graphics technology including ultra-realistic water scenes with river rapids, lifelike bubbles and book pages that literally "wilt" when turning the wet paper.

"There were many technical challenges in the making of 'Ratatouille,'" said Greg Brandeau, senior vice president of technology, Pixar Animation Studios. "Intel's advanced computing capabilities helped Pixar bring 'Ratatouille' to life faster than ever, delivering a 30 percent performance improvement in the computer-generated animation and visual effects rendering software. Faster rendering gives lighting designers more time to create even more realistic images like an animated Eiffel Tower or Remy the rat's winning smile."

"We are absolutely thrilled to be developing our first promotion with Intel," said Cherise McVicar, senior vice president of national promotions and mobile marketing, Disney. "Featuring their innovative technologies, Intel has been a 'soup to nuts' collaborator, supporting both the making of and the marketing of Disney-Pixar's Ratatouille. We're grateful to collaborate with such an iconic technology company in both advertising and engaging consumers."

"Movie magic starts with powerful behind the scenes performances," said CJ Bruno, Intel's director of Americas Marketing. "The same Intel Core technology that gave Disney-Pixar the performance and flexibility needed to create incredible realism in 'Ratatouille,' can

also be used by consumers to watch, produce and share their own movie magic at home and on-the-go."

#### **Intel and the Entertainment Industry**

Over the past decade, Intel has worked closely with Hollywood studios to radically increase the computing capabilities and solutions available to filmmakers, increasing the realism and quality of films. Intel technology – from servers to workstations to software enabling – has played a critical role in creating incredible realism in animated films including "Monsters Inc.," "Finding Nemo," "The Incredibles," and "Cars."

In addition to being a key technology ingredient in the film and helping increase overall studio productivity while reducing production costs, Intel is a major promotional sponsor for "Ratatouille." The campaign, featuring Remy multi-tasking and entertaining with Intel® Core™ 2 Duo processors, includes retail promotions with Toshiba\* and Circuit City\*, special event sponsorships such as Disney's "Big Cheese Tour," [www.thebestingredients.com](http://www.thebestingredients.com) Web site and advertising.

#### **About Intel**

Intel (NASDAQ: INTC), the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at [www.intel.com/pressroom](http://www.intel.com/pressroom) and [blogs.intel.com](http://blogs.intel.com).

#### **About Pixar Animation Studios**

Pixar Animation Studios, a wholly-owned subsidiary of The Walt Disney Company\*, is an Academy Award\*-winning film studio with world-renowned technical, creative and production capabilities in the art of computer animation. Creator of some of the most successful and beloved animated films of all time, including "Toy Story," "Finding Nemo," "The Incredibles" and "Cars," the Northern California studio has won 20 Academy Awards\*; and its seven films have grossed more than \$3.7 billion at the worldwide box office to date. The next scheduled film release from Disney-Pixar is "WALL•E" (June 27, 2008).

Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.


\* Other brands may be claimed as the property of others.

[Back to Top](#)

[Site Map](#) [RSS](#) [Jobs](#) [Investor Relations](#) [Press Room](#) [Contact Us](#)

[Terms of Use](#) [\\*Trademarks](#) [Social Media Guidelines](#) [Privacy](#) ©Intel Corporation

# EXHIBIT 2



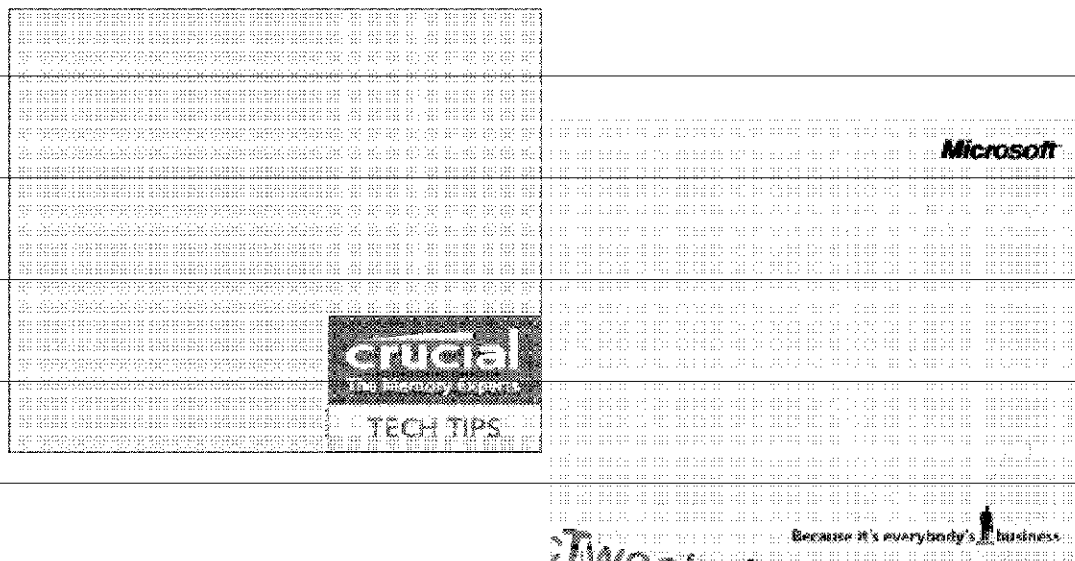
# Intel Quietly Acquires Graphics Software Firm

Intel Takes Over Graphics Firm Neoptica

by **Anton Shilov**

11/26/2007 | 02:01 PM

Intel Corp. has quietly acquired graphics software company called Neoptica. The decision once again outlines importance that graphics technologies represent for the world's largest maker of x86 central processing units (CPUs). Unfortunately, it is unclear what will the company's do as a part of Intel.



"Neoptica has been acquired by Intel, where we are working on a number of exciting efforts in interactive computer graphics," a short statement from Neoptica's [web-site](#) reads.

Previously, Neoptica developed software tools for development of advanced "interactive graphics". Based in San Francisco, California, the company was founded by ex-Nvidia and ex-Sony Computer Entertainment Inc. employees who used to work on software development positions. Besides, some of Neoptica's employees were involved into GPGPU (general purpose computing on graphics processing units) projects.

Presently Intel commands the largest share on the market of graphics adapters with its chipsets with built-in graphics cores. However, it is rumoured that Intel is working on high-performance discrete graphics processing units (GPUs), which is why the company is hiring additional personnel experienced in development of graphics hardware and software.

Intel has not yet released any statements regarding the Neoptica take over.



# EXHIBIT 3

# Larrabee: A Many-Core x86 Architecture for Visual Computing

Larry Seiler<sup>1</sup>, Doug Carmean<sup>1</sup>, Eric Sprangle<sup>1</sup>, Tom Forsyth<sup>1</sup>, Michael Abrash<sup>2</sup>,  
Pradeep Dubey<sup>1</sup>, Stephen Junkins<sup>1</sup>, Adam Lake<sup>1</sup>, Jeremy Sugerman<sup>3</sup>,  
Robert Cavin<sup>1</sup>, Roger Espasa<sup>1</sup>, Ed Grochowski<sup>1</sup>, Toni Juan<sup>1</sup>, and Pat Hanrahan<sup>3</sup>

### Abstract

This paper presents a many-core visual computing architecture code named Larrabee, a new software rendering pipeline, a many-core programming model, and performance analysis for several applications. Larrabee uses multiple in-order x86 CPU cores that are augmented by a wide vector processor unit, as well as some fixed function logic blocks. This provides dramatically higher performance per watt and per unit of area than out-of-order CPUs on highly parallel workloads. It also greatly increases the flexibility and programmability of the architecture as compared to standard GPUs. A coherent on-die 2<sup>nd</sup> level cache allows efficient inter-processor communication and high-bandwidth local data access by CPU cores. Task scheduling is performed entirely with software in Larrabee, rather than in fixed function logic. The customizable software graphics rendering pipeline for this architecture uses binning in order to reduce required memory bandwidth, minimize lock contention, and increase opportunities for parallelism relative to standard GPUs. The Larrabee native programming model supports a variety of highly parallel applications that use irregular data structures. Performance analysis on those applications demonstrates Larrabee's potential for a broad range of parallel computation.

CCS: I.3.1 [Computer Graphics]: Hardware Architecture--Graphics Processors, Parallel Processing; I.3.3 [Computer Graphics]: Picture/Image Generation--Display Algorithms; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism--Color, shading, shadowing, and texture

**Keywords:** graphics architecture, many-core computing, real-time graphics, software rendering, throughput computing, visual computing, parallel processing, SIMD, GPGPU.

### 1. Introduction

Modern GPUs are increasingly programmable in order to support advanced graphics algorithms and other parallel applications.

<sup>1</sup> Intel® Corporation: larry.seiler, doug.carmean, eric.sprangle, tom.forsyth, pradeep.dubey, stephen.junkins, adam.lake, robert.d.cavin, roger.espada, edward.grochowski & toni.juan @intel.com

<sup>2</sup> RAD Game Tools: mikea@radgametools.com

<sup>3</sup> Stanford University: yoel & hanrahan @cs.stanford.edu

#### ACM Reference Format

Seiler, L., Carmean, D., Sprangle, E., Forsyth, T., Abrash, M., Dubey, P., Junkins, S., Lake, A., Sugerman, J., Cavin, R., Espasa, R., Grochowski, E., Juan, T., Hanrahan, P. 2008. Larrabee: A Many-Core x86 Architecture for Visual Computing. *ACM Trans. Graph.* 27, 3, Article 18 (August 2008), 15 pages. DOI = 10.1145/1360812.1360817 <http://doi.acm.org/10.1145/1360812.1360817>.

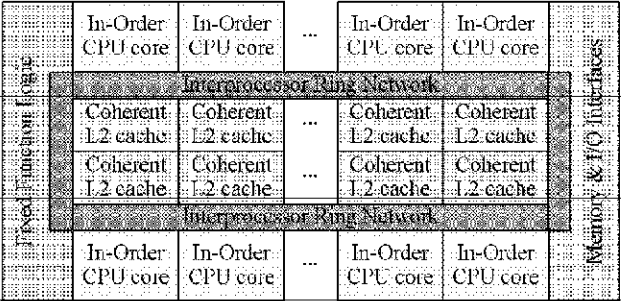
#### Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).  
© 2008 ACM 0730-0301/2008/03-ART18 \$5.00 DOI 10.1145/1360812.1360817  
<http://doi.acm.org/10.1145/1360812.1360817>

However, general purpose programmability of the graphics pipeline is restricted by limitations on the memory model and by fixed function blocks that schedule the parallel threads of execution. For example, pixel processing order is controlled by the rasterization logic and other dedicated scheduling logic.

This paper describes a highly parallel architecture that makes the rendering pipeline completely programmable. The Larrabee architecture is based on in-order CPU cores that run an extended version of the x86 instruction set, including wide vector processing operations and some specialized scalar instructions. Figure 1 shows a schematic illustration of the architecture. The cores each access their own subset of a coherent L2 cache to provide high-bandwidth L2 cache access from each core and to simplify data sharing and synchronization.

Larrabee is more flexible than current GPUs. Its CPU-like x86-based architecture supports subroutines and page faulting. Some operations that GPUs traditionally perform with fixed function logic, such as rasterization and post-shader blending, are performed entirely in software in Larrabee. Like GPUs, Larrabee uses fixed function logic for texture filtering, but the cores assist the fixed function logic, e.g. by supporting page faults.



**Figure 1:** Schematic of the Larrabee many-core architecture: The number of CPU cores and the number and type of co-processors and I/O blocks are implementation-dependent, as are the positions of the CPU and non-CPU blocks on the chip.

This paper also describes a software rendering pipeline that runs efficiently on this architecture. It uses binning to increase parallelism and reduce memory bandwidth, while avoiding the problems of some previous tile-based architectures. Implementing the renderer in software allows existing features to be optimized based on workload and allows new features to be added. For example, programmable blending and order-independent transparency fit easily into the Larrabee software pipeline.

Finally, this paper describes a programming model that supports more general parallel applications, such as image processing, physical simulation, and medical & financial analytics. Larrabee's support for irregular data structures and its scatter-gather capability make it suitable for these throughput applications as demonstrated by our scalability and performance analysis.

2. Previous Work

Recent years have seen the evolution of commodity graphics hardware from fixed function units toward an increasingly programmable graphics pipeline, offering greater flexibility and performance to graphics programmers [Blythe 2006]. Real-time graphics APIs evolved to expose this functionality with high level shading languages such as Cg [Mark et al. 2003], HLSL [Microsoft 2007] and GLSL [Kessenich et al. 2006]. Additionally, a number of arithmetically intensive workloads perform very well on GPU-like architectures [GPGPU 2007; Owens et al. 2007].

2.1 PC Graphics Processor Architectures

Multi-processing graphics hardware has been around for decades. The key ideas behind these architectures are described by Fuchs et al. [1989], Molnar et al. [1992], Foley et al. [1996], and Stoll et al. [2001]. The motivation has always been to leverage the data parallel nature of rendering to gain maximum performance and visual fidelity. Early architectures were complex systems with multiple boards and many specialized chips. The modern graphics architecture is a single chip that fits into the form factor of a PC or other graphics platform [Kelley et al. 1992; Kelley et al. 1994; Torborg & Kajiya 1996]. Recent architectures implement the Microsoft DirectX<sup>®</sup> 10 API, including the Nvidia GeForce<sup>®</sup> 8 [Nvidia 2008] and the ATI Radeon<sup>®</sup> HD 3800 series [AMD 2008].

Figure 2 shows a simplified version of the DirectX 10 pipeline [Blythe 2006]. The programmable OpenGL pipeline is following a similar architectural direction [Rost 2004]. In early implementations, each stage required separate programmable units, but the most recent architectures use a unified shader model. Rasterization and texture filter operations are still largely fixed function in modern GPUs, so changes to the supported features require a new chip design, as well as a new API version.




Figure 2: Simplified DirectX10 Pipeline: Yellow components are programmable by the user, green are fixed function. Memory access, stream output, and texture filtering stages are omitted.

Recent research brings the computational capabilities of commodity graphics hardware to bear on a variety of problems including collision detection, financial modeling, and signal processing [Owens et al. 2007]. There are several aspects of GPU architectures that lend themselves to these workloads. Performance increases are most dramatic when the compute to bandwidth ratio is high. That occurs when the application primarily uses regular data structures and requires many arithmetic operations for each data item being processed.

2.2 Taxonomy of Rendering Methods

Molnar et al. [1994] classified graphics architectures by the primary stage in the graphics pipeline where sorting for parallelism occurs. Eldridge [2001] expanded this taxonomy and performed a detailed comparison of parallel rendering systems. Both use *sort-middle* to refer to sorting after geometry processing but before rasterization. Molnar et al. use *sort-last sparse* and

Eldridge uses *sort-last fragment* to refer to sorting before fragment processing, that is, before depth test and alpha blending. Since their work was done before the introduction of pixel shaders, no clear distinction is provided for sorting after computing coverage but before executing the pixel shader. In this paper we treat this alternative as a variant of *sort-middle*.

Current GPUs often use Eldridge’s sort-last fragment. This allows pixel shader results to be sorted into a small number of screen-aligned regions just before depth testing and blending. Often these screen aligned regions are associated with individual memory controllers, to allow more efficient memory accesses. Sort-last fragment allows immediate mode rendering with relatively short FIFO buffers for sorting the data. However, if a pixel is accessed multiple times at widely separated intervals, it typically must be read from and written to memory multiple times.

Sort middle algorithms have been called binning, tiling, chunking, bucket, and zone rendering [Hsieh et al. 2001; Chen et al. 1998]. This method processes vertices to produce screen coordinates for primitives, which are sorted into bins based on their location on the screen. Each bin is associated with a tile region on the screen, which can be sized to fit into on-chip cache so that only one access to memory is needed per pixel regardless of the depth complexity. One problem is that primitives that overlap multiple tiles must be stored in multiple bins, which increases the memory bandwidth needed for them. Molnar et al. [1994], Chen et al. [1998], and Eldridge [2001] concluded that the impact of overlap is limited, especially when primitives are small compared to region size, which occurs as the triangle count of a scene increases. The PowerVR<sup>®</sup> MBX and SGX series [PowerVR 2008], the Intel<sup>®</sup> Graphics Media Accelerator 900 Series [Lake 2005], the ARM Mali<sup>®</sup> [Stevens 2006], and Talisman [Torborg & Kajiya 1996] have been generally classified as sort middle architectures.

2.3 General Purpose CPU Architectures

In 1995, Intel introduced the Pentium<sup>®</sup> Pro processor, which used out-of-order 3-wide instruction execution in response to the demand for increasing single-stream performance [Gwennap 1995]. Out-of-order architectures identify independent instruction streams that can be executed in parallel. The logic to identify these instructions consumes die area as well as power. Later CPU generations used even more elaborate techniques to increase single-stream performance at ever increasing penalties in area and power relative to performance. It has been observed that, within the same process technology, a new microprocessor design with 1.5x to 1.7x the performance consumes 2x to 3x the die area [Pollack 1999] and 2x to 2.5x the power [Grochowski et al. 2004].

For highly parallel algorithms, more performance can be gained by packing multiple cores onto the die instead of increasing single stream performance. The IBM Cell<sup>®</sup> was designed with these high throughput workloads in mind [Pham et al. 2005]. Cell includes a single Power<sup>®</sup> Processor core, its L2 cache, and a set of high throughput cores. These cores each contain a local memory store that is incoherent with the rest of the memory system. The local store has a guaranteed latency for data delivery, which allows a simpler execution pipeline than a system with a coherent cache hierarchy. It requires the user to manually manage the data contents through software-programmed DMA operations.

Intel, Intel Core, Pentium and Xeon are trademarks of Intel Corporation in the U.S. and other countries.  
\* Other names & brands may be claimed as the property of others.

CPUs can also use multi-threading to gain parallelism. Niagara is a multi-core general-purpose microprocessor [Kongetira et al. 2005] featuring eight in-order cores, each capable of executing four simultaneous threads, and a shared cache. But given its focus on commercial server workloads, Niagara lacks architectural elements critical for visual computing, such as SIMD floating-point execution, scatter-gather, or fixed function texture support.

3. Larrabee Hardware Architecture

Figure 1 above shows a block diagram of the basic Larrabee architecture. Larrabee is designed around multiple instantiations of an in-order CPU core that is augmented with a wide vector processor (VPU). Cores communicate through a high-bandwidth interconnect network with some fixed function logic, memory I/O interfaces, and other necessary I/O logic, depending on the exact application. For example, an implementation of Larrabee as a stand-alone GPU would typically include a PCIe bus.

The data in Table 1 motivates Larrabee’s use of in-order cores with wide VPUs. The middle column shows the peak performance of a modern out-of-order CPU, the Intel® Core™2 Duo processor. The right-hand column shows a test CPU design based on the Pentium® processor, which was introduced in 1992 and used dual-issue in-order instruction execution [Alpert 1993]. The Pentium processor core was modified to support four threads and a 16-wide VPU. The final two rows specify the number of non-vector instructions that can be issued per clock by one CPU and the total number of vector operations that can be issued per clock. The two configurations use roughly the same area and power.

# CPU cores:	2 out-of-order	10 in-order
Instruction issue:	4 per clock	2 per clock
VPU per core:	4-wide SSE	16-wide
L2 cache size:	4 MB	4 MB
Single-stream:	4 per clock	2 per clock
Vector throughput:	8 per clock	160 per clock

Table 1: Out-of-order vs. in-order CPU comparison: designing the processor for increased throughput can result in 1/2 the peak single-stream performance, but 20x the peak vector throughput with roughly the same area and power. This difference is 40x in FLOPS, since the wide VPU supports fused multiply-add but SSE doesn’t. These in-order cores are not Larrabee, but are similar.

The test design in Table 1 is not identical to Larrabee. To provide a more direct comparison, the in-order core test design uses the same process and clock rate as the out-of-order cores and includes no fixed function graphics logic. This comparison motivates design decisions for Larrabee since it shows that a wide VPU with a simple in-order core allows CPUs to reach a dramatically higher computational density for parallel applications.

Sections 3.1 to 3.5 below describe the key features of the Larrabee architecture: the CPU core, the scalar unit and cache control instructions, the vector processor, the interprocessor ring network, and the choices for what is implemented in fixed function logic.

3.1 Larrabee Core and Caches

Figure 3 shows a schematic of a single Larrabee CPU core, plus its connection to the on-die interconnect network and the core’s local subset of the L2 cache. The instruction decoder supports the standard Pentium processor x86 instruction set, with the addition of new instructions that are described in Sections 3.2 and 3.3. To

simplify the design the scalar and vector units use separate register sets. Data transferred between them is written to memory and then read back in from the L1 cache.

Larrabee’s L1 cache allows low-latency accesses to cache memory into the scalar and vector units. Together with Larrabee’s load-op VPU instructions, this means that the L1 cache can be treated somewhat like an extended register file. This significantly improves the performance of many algorithms, especially with the cache control instructions described Section 3.2. The single-threaded Pentium processor provided an 8KB Icache and 8KB Dcache. We specify a 32KB Icache and 32KB Dcache to support four execution threads per CPU core.




Figure 3: Larrabee CPU core and associated system blocks: the CPU is derived from the Pentium processor in-order design, plus 64-bit instructions, multi-threading and a wide VPU. Each core has fast access to its 256KB local subset of a coherent 2<sup>nd</sup> level cache. L1 cache sizes are 32KB for Icache and 32KB for Dcache. Ring network accesses pass through the L2 cache for coherency.

Larrabee’s global 2<sup>nd</sup> level (L2) cache is divided into separate local subsets, one per CPU core. Each CPU has a fast direct access path to its own local subset of the L2 cache. Data read by a CPU core is stored in its L2 cache subset and can be accessed quickly, in parallel with other CPUs accessing their own local L2 cache subsets. Data written by a CPU core is stored in its own L2 cache subset and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data, as described in Section 3.4. We specify 256KB for each L2 cache subset. This supports large tile sizes for software rendering, as described in Section 4.1.

3.2 Scalar Unit and Cache Control Instructions

Larrabee’s scalar pipeline is derived from the dual-issue Pentium processor, which uses a short, inexpensive execution pipeline. Larrabee provides modern additions such as multi-threading, 64-bit extensions, and sophisticated prefetching. The cores support the full Pentium processor x86 instruction set so they can run existing code including operating system kernels and applications. Larrabee adds new scalar instructions such as bit count and bit scan, which finds the next set bit within a register.

Larrabee also adds new instructions and instruction modes for explicit cache control. Examples include instructions to prefetch data into the L1 or L2 caches and instruction modes to reduce the priority of a cache line. For example, streaming data typically sweeps existing data out of a cache. Larrabee is able to mark each streaming cache line for early eviction after it is accessed. These cache control instructions also allow the L2 cache to be used similarly to a scratchpad memory, while remaining fully coherent.

Within a single core, synchronizing access to shared memory by multiple threads is inexpensive. The threads on a single core share the same local L1 cache, so a single atomic semaphore read within the L1 cache is sufficient. Synchronizing access between

multiple cores is more expensive, since it requires inter-processor locks. This is a well-known problem in multi-processor design.

Multi-issue CPU cores often lose performance due to the difficulty of finding instructions that can execute together. Larrabee's dual-issue decoder has a high multi-issue rate in code that we've tested. The pairing rules for the primary and secondary instruction pipes are deterministic, which allows compilers to perform offline analysis with a wider scope than a runtime out-of-order instruction picker can. All instructions can issue on the primary pipeline, which minimizes the combinatorial problems for a compiler. The secondary pipeline can execute a large subset of the scalar x86 instruction set, including loads, stores, simple ALU operations, branches, cache manipulation instructions, and vector stores. Because the secondary pipeline is relatively small and cheap, the area and power wasted by failing to dual-issue on every cycle is small. In our analysis, it is relatively easy for compilers to schedule dual-issue instructions.

Finally, Larrabee supports four threads of execution, with separate register sets per thread. Switching threads covers cases where the compiler is unable to schedule code without stalls. Switching threads also covers part of the latency to load from the L2 cache to the L1 cache, for those cases when data cannot be prefetched into the L1 cache in advance. Cache use is more effective when multiple threads running on the same core use the same dataset, e.g. rendering triangles to the same file.

3.3 Vector Processing Unit

Larrabee gains its computational density from the 16-wide vector processing unit (VPU), which executes integer, single-precision float, and double-precision float instructions. The VPU and its registers are approximately one third the area of the CPU core but provide most of the integer and floating point performance. Figure 4 shows a block diagram of the VPU with the L1 cache.




Figure 4: Vector unit block diagram: the VPU supports 3-operand instructions. It supports swizzling the register inputs and numeric conversion and replication on the memory input. Mask registers allow predicating the resulting vector writes.

We chose a 16-wide VPU as a tradeoff between increased computational density and the difficulty of obtaining high utilization for wider VPUs. Early analysis suggested 88% utilization for typical pixel shader workloads if 16 lanes process 16 separate pixels one component at a time, that is, with separate instructions to process red, green, etc., for 16 pixels at a time, instead of processing multiple color channels at once. The Nvidia GeForce 8 operates in a similar fashion, organizing its scalar SIMD processors in groups of 32 that execute the same instruction [Nickolls et al. 2008]. The main difference is that in Larrabee the loop control, cache management, and other such operations are code that runs in parallel with the VPU, instead of being implemented as fixed function logic.

Larrabee VPU instructions allow up to three source operands, one of which can come directly from the L1 cache. If the data has

been prefetched into the cache, as described in Section 3.2, then the L1 cache is in effect an extended register file. 8-bit unorm, 8-bit uint, 16-bit sint and 16-bit float data can be read from the cache and converted to 32-bit floats or 32-bit integers with no loss of performance. This significantly increases the amount of data that can be stored in the caches and also reduces the need for separate data conversion instructions.

The next stage is to align the data from registers and memory with the processing lanes in the VPU. Register data can be swizzled in a variety of ways, e.g. to support matrix multiplication. Data from memory can be replicated across the VPU lanes. This is a common operation in both graphics and non-graphics parallel data processing, which significantly increases the cache efficiency.

The VPU supports a wide variety of instructions on both integer and floating point data types. The instruction set provides the standard arithmetic operations, including fused multiply-add, and the standard logical operations, including instructions to extract non-byte-aligned fields from pixels. These are all load-op instructions, which read from registers or memory and write the result to a vector register. Additional load and store instructions support a wider variety of conversions between floating point values and the less common or more complex data formats found on most GPUs. Using separate instructions for these formats saves significant area and power at a small performance cost.

The VPU instruction set also includes gather and scatter support, that is, loads and stores from non-contiguous addresses. Instead of loading a 16-wide vector from a single address, 16 elements are loaded from or stored to up to 16 different addresses that are specified in another vector register. This allows 16 shader instances to be run in parallel, each of which appears to run serially, even when performing array accesses with computed indices. The speed of gather/scatter is limited by the cache, which typically only accesses one cache line per cycle. However, many workloads have highly coherent access patterns, and therefore take much less than 16 cycles to execute.

Finally, Larrabee VPU instructions can be predicated by a mask register, which has one bit per vector lane. The mask controls which parts of a vector register or memory location are written and which are left untouched. For example, a scalar if-then-else control structure can be mapped onto the VPU by using an instruction to set a mask register based on a comparison, and then executing both if and else clauses with opposite polarities of the mask register controlling whether to write results. Clauses can be skipped entirely if the mask register is all zeros or all ones. This reduces branch misprediction penalties for small clauses and gives the compiler's instruction scheduler greater freedom.

The VPU also uses these masks for packed load and store instructions, which access enabled elements from sequential locations in memory. This enables the programmer to bundle sparse strands of execution satisfying complex branch conditions into a format more efficient for vector computation.

3.4 Inter-Processor Ring Network

Larrabee uses a bi-directional ring network to allow agents such as CPU cores, L2 caches and other logic blocks to communicate with each other within the chip. When scaling to more than 16 cores, we use multiple short linked rings.

Each ring data-path is 512-bits wide per direction. All the routing decisions are made before injecting messages into the network. For example, each agent can accept a message from one direction



on even clocks and from the other direction on odd clocks. This simplifies the routing logic and means that no storage is required in the routers once the message is in the network. The result is high bandwidth with minimal contention at a very low cost.

Larrabee's L2 cache is designed to provide each core with high bandwidth access to memory addresses that are not written by other cores, and therefore are stored in the core's local L2 subset. Each core can access its own subset of the L2 cache in parallel, without communicating with other cores. However, before allocating a new line in the L2 cache, the ring network is used to check for data sharing, in order to maintain data coherency.

The inter-processor network also provides a path for the L2 caches to access memory. A typical high-end implementation would include multiple memory interfaces of standard design, spread around the inter-processor network to reduce congestion. Latency around the on-die network increases memory access times, but the extra ring latency is typically very small compared to the latency of DRAM access.

Finally, the on-die inter-processor network allows fixed function logic agents to be accessed by the CPU cores and in turn to access L2 caches and memory. As with memory controllers, these would typically be spread around the ring network to reduce congestion.

3.5 Fixed Function Logic

Modern GPUs contain fixed function logic for a variety of graphics tasks, including texture filtering, display processing, post-shader alpha blending, rasterization, and interpolation. In this paper, *rasterization* refers solely to finding the coverage of a primitive, and *interpolation* refers to finding the values of parameters at covered sample positions in the primitive. Fixed function logic typically requires FIFOs for load balancing. It can be difficult to properly size these logic blocks and their FIFOs to avoid both wasted area and performance bottlenecks.

Larrabee uses software in place of fixed function logic when a software implementation provides sufficient performance. In particular, Larrabee does not include fixed function logic for rasterization, interpolation, or post-shader alpha blending. This allows Larrabee to add new features and optimizations, as well as allowing these tasks to be implemented in different places in the rendering pipeline, depending what is most efficient for a particular application. Implementing them in software also allows Larrabee to allocate to each the performance it requires, instead of designing hardware to meet peak performance requirements. Sections 4.4 and 4.5 describe the software algorithms used and Section 5.5 shows the percentage of processing time required by these operations for three game workloads.

Larrabee includes texture filter logic because this operation cannot be efficiently performed in software on the cores. Our analysis shows that software texture filtering on our cores would take 12x to 40x longer than our fixed function logic, depending on whether decompression is required. There are four basic reasons:

- Texture filtering still most commonly uses 8-bit color components, which can be filtered more efficiently in dedicated logic than in the 32-bit wide VPU lanes.
- Efficiently selecting unaligned 2x2 quads to filter requires a specialized kind of pipelined gather logic.
- Loading texture data into the VPU for filtering requires an impractical amount of register file bandwidth.
- On-the-fly texture decompression is dramatically more efficient in dedicated hardware than in CPU code.

The Larrabee texture filter logic is internally quite similar to typical GPU texture logic. It provides 32KB of texture cache per core and supports all the usual operations, such as DirectX 10 compressed texture formats, mipmapping, anisotropic filtering, etc. Cores pass commands to the texture units through the L2 cache and receive results the same way. The texture units perform virtual to physical page translation and report any page misses to the core, which retries the texture filter command after the page is in memory. Larrabee can also perform texture operations directly on the cores when the performance is fast enough in software.

4. Larrabee Software Renderer

The key issue for achieving high performance for any parallel rendering algorithm is to divide the rendering task into many subtasks that can be load balanced and executed in parallel with very few synchronization points. Larrabee allows more options for parallelism than typical GPUs due to its flexible memory model and software-controlled scheduling.

This section describes a sort-middle software renderer designed for the Larrabee architecture that uses binning for load balancing. Section 5 provides performance studies for this software renderer.

4.1 Stages of Software Rendering

For simplicity, first we will consider rendering to a single set of render targets, such as a pixel buffer and a depth/stencil buffer. These render targets and the rendering commands that modify them are together called an *RTset*. Section 4.2 discusses more complex cases involving multiple RTsets.

The rendering commands for an RTset are typically specified by graphics APIs as a series of rendering state changes, followed by a batch of triangles rendered using that current device state. Rather than use the concept of a *current* state internally, the Larrabee renderer captures the rendering state in a single fully-specified structure. It then groups the batches of triangles and tags each batch with the state it uses. This batch of triangles and the state it uses is called a *primitive set* or *PrimSet*. This is roughly equivalent to the DirectX DrawPrimitive call, although there is not an exact 1:1 correspondence between the two.

Figure 5 shows the broad structure for rendering the PrimSets of a single RTset. The surface being rendered is split into tiles of pixels. Each tile has a bin that will be filled with the triangles from a PrimSet that intersect that tile. The set of bins for the whole RTset is called a *bin set*. The terms *tile* and *bin* are sometimes used interchangeably. The distinction in this paper is that a *tile* is the actual pixel data, while the *bin* is the set of primitives that affect that tile. In the same way that each tile has a bin, each RTset (set of render target tiles and associated PrimSets) has a single bin set (set of bins that contain the primitives).




Figure 5: Larrabee Software Renderer Structure: Multiple sets of primitives (PrimSets) can be processed in parallel to fill per-tile bins, which are later processed in parallel to render screen tiles.

18:6 • L. Seiler et al.

Tile size is chosen so that the target surfaces in the RTset for that tile will all fit in a core's L2 cache. Thus an RTset with many color channels, or with large high-precision data formats, will use a smaller tile size than one with fewer or low-precision channels. To simplify the code, tiles are usually square and a power-of-two in size, typically ranging in size from 32x32 to 128x128. An application with 32-bit depth and 32-bit color can use a 128x128 tile and fill only half of the core's 256KB L2 cache subset.

As long as a tile fits within the L2 cache, rendering speed does not change substantially for different tile sizes. The main impact of using smaller tiles is that some triangles in the scene will hit more than one tile and require processing in each of those tiles – this is termed *bin spread*. Smaller tiles increase bin spread, but it is not a large increase. Typically we see bin spread of less than 5% in modern workloads. That is, the number of triangles processed across the system is less than 5% higher than the number for a single large bin covering the entire render target.

There are two phases to the processing. In the front-end, each PrimSet is given a sequence ID to identify where in the rendering stream it was submitted. This is used by the back-end to ensure correct ordering, as discussed below. The PrimSet is then assigned to a single core, which performs vertex shading, tessellation, geometry shading, culling and clipping to produce triangles (or other primitives). The core then rasterizes each triangle to determine which tiles it touches and which samples it covers within each of those tiles. The result is a series of X,Y coordinates and sample coverage masks for each triangle. This data is stored in the bins along with indices that reference the vertex data.

Once all front-end processing for the RTset has finished and every triangle has been added to the bin for each tile that it touched, back-end processing is performed. Here, each tile is assigned to a single core, which shades each triangle from the associated bin, including requesting texture sampling from the co-processors. The back-end also performs depth, stencil and blending operations.


Unlike some other tile-based rendering methods, there is no attempt at perfect occlusion culling before shading, reordering of shading, or any other non-standard rendering methods. When taking commands from a DirectX or OpenGL command stream, rendering for a single tile is performed in the order in which the commands are submitted. Using a conventional rendering pipeline within each tile avoids surprises in either functionality or performance and works consistently well across a broad spectrum of existing applications.

## 4.2 Render Target Dependency Analysis

A single frame consists of a sequence of rendering commands, each sent to a set of rendering surfaces. Modern applications may use multiple pixel targets at once, and may change targets frequently during a single frame in order to render effects such as reflections and shadow maps.

To handle different sets of render targets within a single frame, Larrabee's software renderer starts by creating a graph where each node corresponds to an *RTset*, as defined in Section 4.1. Each node is then assigned the PrimSets that modify that node's set of render targets. When an RTset uses a render target (e.g. a texture) that is used by subsequent rendering operations to a different target, a dependency is set up between the two RTsets. For example, in shadow mapping, the main RTset for a scene (the back buffer and depth/stencil buffer) has a dependency on the RTset for each of the shadow maps used.

Once the dependency graph is created, the nodes can be selected for rendering in any order that satisfies the dependencies. Figure 6 shows a dependency graph for two frames of a scene that requires rendering two shadow maps. For simplicity, the shadow maps for frame 1 are not shown. Frame 2 of the scene cannot be rendered until after frame 2's shadow maps are rendered. Since each frame in this simple example uses the same memory for the back buffer and depth buffer, frame 2 also cannot be rendered until frame 1's scene is rendered and copied to the front buffer (the dotted line dependency). However, rendering the frame 2 shadow maps can overlap with frame 1 rendering, since there are no dependencies. Using a different back buffer for frame 2 would remove the dotted line dependency. This substitution can be done automatically.



**Figure 6:** RTset dependency graph: PrimSets are assigned to an RTset node based on the surfaces (render targets) that they modify. The dependencies ensure that a surface is not used until the PrimSets that modify it have been rendered and is not modified until the PrimSets that use it have been rendered.

Note that the PrimSets associated with an RTset can be divided into multiple subsets whenever required, so long as ordering is maintained. An RTset can be split if it is too large to be efficiently processed as one unit, e.g. to provide finer scheduling granularity. Not all Larrabee cores need to process PrimSets from the same RTset at the same time. This ability to arbitrarily split and schedule RTsets avoids the limitations of some previous tiling architectures [Lake 2005].


## 4.3 Front-End Vertex and Geometry Processing

Since graphics rendering commands modify state, the order of execution matters. GPUs process these commands sequentially, so that the commands are started in order and finished in order. When operations within a rendering command are parallelized over the inputs, the outputs must be put back in order. Geometry shaders, where the number of outputs is variable, require particularly large FIFOs to maintain order and minimize stalls.

Larrabee allows front-end processing of multiple PrimSets in parallel. A control processor decides which PrimSets to render at any particular time, according to the dependencies in the RTset graph, and adds those PrimSets to an active list. The Larrabee cores doing front-end work constantly take PrimSets from this active list. Each core works on its own PrimSet independently. When the core is finished, it takes the next from the active list. Each core uses its own subset of the bin for each tile, which eliminates lock contention with the other front-end cores. The PrimSet's sequence ID is written into the bins so that the back-end can restore the original order by always reading primitives from the sub-bin with the smallest sequence ID.

Figure 7 shows the processing stages within a single front-end core. The first step identifies the vertices that form each primitive. This can be complex due to index buffers that allow arbitrary mappings of vertices in a vertex buffer to primitives, e.g. to efficiently store meshes. Next, the required vertices are transformed by running the vertex shader on them if they haven't already been transformed. Transformed vertices are streamed out

to main memory. Values other than the position data are actively evicted from the L2 cache to avoid pollution, as they are not needed again until interpolant setup in the back end. After this, the geometry shader is run, followed by frustum and back-face culling, then clipping.



**Figure 7: Front-End Rendering Sequence for a PrimSet:** the renderer shades vertices when required for a primitive, then puts final primitives into the bins whose tiles the primitive intersects.

We describe a version of the algorithm that computes coverage information in the front-end and puts it into the bins. This ensures good load balancing, even if a small number of bins contain a large number of triangles. Rasterization can occur in either the front-end or the back-end, or can be split between them, since Larrabee uses software rasterization, as described in Section 4.4.

#### 4.4 Software Rasterization and Interpolation

Unlike modern GPUs, Larrabee does not use dedicated logic for rasterization and parameter interpolation. In this paper, *rasterization* refers to finding the coverage of a primitive and *interpolation* refers to finding the values of parameters at covered sample positions. The figures in Section 5.4 show that these operations do not take a significant fraction of the rendering workload, so using software is justifiable. This section describes the algorithms that we use and why software implementations are preferable on Larrabee.

The justification for performing interpolation in software is relatively simple. In older graphics APIs, interpolation produced fixed point numbers, much like the current state for the most common texture filtering operations. In modern graphics APIs such as DirectX 10, the required result is a 32-bit float. Therefore it is efficient to re-use the existing VPU for interpolation.

Rasterization is unquestionably more efficient in dedicated logic than in software when running at peak rates, but using dedicated logic has drawbacks for Larrabee. In a modern GPU, the rasterizer is a fine-grain serialization point: all primitives are put back in order before rasterization. Scaling the renderer over large numbers of cores requires eliminating all but the most coarse-grained serialization points. The rasterizer could be designed to allow multiple cores to send it primitives out of order, but this would impose a significant communication expense and would require software to manage contention for the rasterizer resource. A software rasterizer avoids these costs. It also allows rasterization to be parallelized over many cores or moved to multiple different places in the rendering pipeline. We can optimize the rasterization code for a particular workload or support alternative rasterization equations for special purposes [Lloyd et al., 2007].

Our algorithm is a highly optimized version of the recursive descent algorithm described by Greene [1996]. The basic idea is to convert clipped triangles to screen space, then compute a half-plane equation for each triangle edge [Pineda 1988]. This lets us determine if a rectangular block is entirely inside the triangle, entirely outside the triangle, or partially covered by the triangle. In the latter case, the algorithm subdivides the block recursively until it is reduced to an individual pixel or sample position.


On Larrabee, the first step uses the triangle's bounding box to find the tiles that the triangle overlaps. In the remaining steps, the VPU computes half-plane equations for 16 blocks at a time. For example, if the tile size is 64x64, the first stage processes 16 16x16 blocks that cover the tile. The find first bit instruction makes it efficient to find fully and partially covered blocks. Detecting fully covered blocks early is important for efficiency. The second stage tests the 16 4x4 sub-blocks of each partially covered 16x16 block. The third stage tests the 16 pixels of each partially covered 4x4 block. This stage can be repeated for multiple sample positions in each pixel. About 70% of the instructions run on the VPU and take advantage of Larrabee's computational density. About 10% of the efficiency of the algorithm comes from special instructions such as find first bit.

#### 4.5 Back-End Pixel Processing

Once the front-end processing for an RTset has completed filling the bins with triangle data, the RTset is put into an active list. The cores doing back-end work constantly take the next available tile from the list and render the triangles in the corresponding bin. This software can use many optimizations that are commonly implemented in fixed function logic in modern GPUs, such as fast clear, hierarchical Z, and early Z tests [Morein 2000]. Hierarchical Z tests can be done in the front-end to reduce the number of primitives placed in the bins.

The back-end code starts by prefetching the render target pixels into the L2 cache. All rendering will then be performed to the L2 cache until there are no more primitives to render for the tile, when it will be written back to memory. As a result, the pixels in the RTset for this tile only need to be read and written once to main memory, regardless of how many overlapping primitives are in the bin. Two important optimizations can also be detected to save substantial memory bandwidth. The read can be eliminated if the first command clears the entire tile. The write can also be eliminated or reduced for depth data that is not required after rendering and for MSAA colors that can be resolved to one color per pixel before writing to memory.

Figure 8 shows a back-end implementation that makes effective use of multiple threads that execute on a single core. A setup thread reads primitives for the tile. Next, the setup thread interpolates per-vertex parameters to find their values at each sample. Finally, the setup thread issues pixels to the work threads in groups of 16 that we call a *quad*. The setup thread uses scoreboarding to ensure that quads are not passed to the work threads until any overlapping pixels have completed processing.



**Figure 8: Back-End Rendering Sequence for a Tile:** one setup thread processes the primitives and assigns them to one of three work threads that do early Z depth tests, pixel shader processing, late Z depth tests, and alpha blending.

The three work threads perform all remaining pixel processing, including pre-shader early Z tests, the pixel shader, regular late Z tests, and post-shader blending. Modern GPUs use dedicated logic for post-shader blending, but Larrabee uses the VPU. This is particularly efficient since many shaders do not use post-shader blending, so that dedicated blending logic can be unused for some



18.8 • L. Seiler et al.

shaders and may limit performance for other shaders. Section 5.4 provides breakdowns of the total processing time devoted to post-shader blending and parameter interpolation.

One remaining issue is texture co-processor accesses, which can have hundreds of clocks of latency. This is hidden by computing multiple qquads on each hardware thread. Each qquad's shader is called a *fiber*. The different fibers on a thread co-operatively switch between themselves without any OS intervention. A fiber switch is performed after each texture read command, and processing passes to the other fibers running on the thread. Fibers execute in a circular queue. The number of fibers is chosen so that by the time control flows back to a fiber, its texture access has had time to execute and the results are ready for processing.

## 5. Renderer Performance Studies

This section describes performance and scaling studies for the Larrabee software renderer described in Section 4. Studies include scalability experiments for software rendering, load balancing studies, bandwidth comparisons of binning to immediate mode renderers, performance on several game workloads, and charts showing the how total processing time is divided among different parts of the software renderer.

### 5.1 Game Workloads and Simulation Method

Performance tests use workloads derived from three well-known games: Gears of War<sup>\*</sup>, F.E.A.R.<sup>\*</sup>, and Half Life<sup>2</sup> Episode 2. Table 2 contains information about the tested frames from each game. Since we are scaling out to large numbers of cores we use a high-end screen size with multisampling when supported.

Half Life 2 ep. 2	F.E.A.R.	Gears of War
1600x1200 4 sample 25 frames (1 in 30) Valve Corp.	1600x1200 4 sample 25 frames (1 in 100) Monolith Productions	1600x1200 1 sample 25 frames (1 in 250) Epic Games Inc

**Table 2:** Workload summary for the three tested games: the frames are widely separated to catch different scene characteristics as the games progress.


We captured the frames by intercepting the DirectX 9 command stream being sent to a conventional graphics card while the game was played at a normal speed, along with the contents of textures and surfaces at the start of the frame. We tested them through a functional model to ensure the algorithms were correct and that the right images were produced. Next, we estimated the cost of each section of code in the functional model, being aggressively pessimistic, and built a rough profile of each frame. We wrote assembly code for the highest-cost sections, ran it through cycle-accurate simulators, fed the clock cycle results back into the functional model, and re-ran the traces. This iterative cycle of refinement was repeated until 90% of the clock cycles executed during a frame had been run through the simulators, giving the overall profiles a high degree of confidence. Texture unit throughput, cache performance and memory bandwidth limitations were all included in the various simulations.

In these studies we measure workload performance in terms of *Larrabee units*. A *Larrabee unit* is defined to be one Larrabee core running at 1 GHz. The clock rate is chosen solely for ease of calculation, since real devices would ship with multiple cores and

a variety of clock rates. Using Larrabee units allows us to compare performance of Larrabee implementations with different numbers of cores running at different clock rates. A single Larrabee unit corresponds to a theoretical peak throughput of 32 GFLOPS, counting fused multiply-add as two operations.

### 5.2 Scalability Studies


The Larrabee software renderer is designed to allow efficient load balancing over a large number of cores. Figure 9 shows the results of testing load balancing for six configurations, each of which scales the memory bandwidth and texture filtering speed relative to the number of cores. This test uses the simulation methodology described in Section 5.1 in combination with a time-based performance model that tracks dependencies and scheduling. This tool is used for multiple graphics products within Intel.



**Figure 9:** Relative Scaling as a Function of Core Count: This shows configurations with 8 to 48 cores, with each game's results plotted relative to the performance of an 8-core system.

The results of the load balancing simulation show a falloff of 7% to 10% from a linear speedup at 48 cores. For these tests, PrimSets are subdivided if they contain more than 1000 primitives, as described in Section 4.2. Additional tests show that F.E.A.R. falls off by only 2% if PrimSets are subdivided into groups of 200 primitives, so code tuning should improve the linearity.

Figure 10 shows the number of Larrabee units required to render sample frames from the three games at 60 frames/second. These results were simulated on a single core with the assumption that performance scales linearly. For Half Life 2 episode 2, roughly 10 Larrabee Units are sufficient to ensure that all frames run at 60 fps or faster. For F.E.A.R. and Gears of War, roughly 25 Larrabee Units suffice.



**Figure 10:** Overall performance: shows the number of Larrabee Units (cores running at 1 GHz) needed to achieve 60fps for of the series of sample frames in each game.

The remaining issue that can limit scalability is software locks. Simulating multiple frames of rendering at such a fine level of detail is extremely costly. However, this software rendering pipeline was explicitly designed to minimize the number of locks

<sup>\*</sup> Other names & brands may be claimed as the property of others.

and other synchronization events. In general, a lock is obtained and released for each of the following events:


- Twice when a bin set is added to the list of work to do, (once for the front-end queue and once for the back-end queue)
- When a PrimSet is processed by a front-end thread or a tile is processed by a back-end thread
- A few low-frequency locks such as resource creation and deletion, buffer modification by the host CPU, and frame flips or presents.

Modern games usually have significantly less than 10,000 locks per frame. The Larrabee ring network provides relatively good performance for low-contention locks of around 100 clocks per lock per core. Together, these numbers are low enough that lock scaling should be fairly linear with the number of cores, given sufficient memory bandwidth.

### 5.3 Binning and Bandwidth Studies

We adopted a binning algorithm primarily to minimize software locks, but it also benefits load balancing and memory bandwidth.


Our algorithm assigns back-end tiles to any core that is ready to process one, without attempting to load balance. In theory this could result in significant load imbalance, though cores are free to start processing the next RTset or switch to front-end processing. Bin imbalance is not a problem in the game workloads we have studied. Figure 11 shows a trace of the back-end bin processing time for 16 frames of Gears of War. Each trace records the processing time for each bin on the screen for a frame, sorted from the fastest to slowest bin, and normalized to 1.0 as the mean bin processing time for that frame. Most bins fall in the range  $\frac{1}{2}x$  to  $2x$  the mean processing time. Few exceed  $3x$  the mean. The other two games produce similar results.



**Figure 11: Bin Balance for Gears of War:** each curve shows the time required to process one frame's bins, in sorted order from fastest to slowest, normalized by the mean bin processing time.

Memory bandwidth is important because the memory subsystem can be one of the more costly and power hungry parts of a GPU, from high end down to low cost designs. It is often a limited resource that can cause bottlenecks if not carefully managed, in part because computational speed scales faster. Our performance studies measure computational speed, unrestricted by memory bandwidth, but it is important to consider how our binning method compares with standard immediate mode rendering algorithms.

Figure 12 compares the total memory bandwidth per frame that we calculated for immediate mode and binned rendering for the three games. The graph presents per-frame data in sorted order from least to most bandwidth for the immediate mode frames. For immediate mode we assume perfect hierarchical depth culling, a 128KB texture cache, and 1MB depth and color caches to represent an ideal implementation. We further assume 2x color and 4x depth compression for single-sampling and 4x color and 8x depth compression for 4-samples per pixel.




**Figure 12: Bandwidth comparison of binning vs. immediate mode per frame:** binning requires bin reads & writes, but eliminates many depth/color accesses that are not detected by hierarchical depth tests. This results in less total bandwidth for binning.

Immediate mode uses more bandwidth for every tested frame: 2.4x to 7x more for F.E.A.R., 1.5x to 2.6x more for Gears of War, and 1.6x to 1.8x more for Half Life 2 episode 2. Notably, binning achieves its greatest improvement when the immediate mode bandwidth is highest, most likely because overdraw forces multiple memory accesses in immediate mode. Even with depth culling and frame buffer compression, the 1MB caches are not large enough to catch most pixel overdraw. High resolutions tend to increase the advantage of binning since they increase the impact of pixel access bandwidth on performance.

### 5.4 Performance Breakdowns


Figure 13 shows the average time spent in each rendering stage for the three games. Pixel shading and interpolant setup is always a major portion of the rendering time, but the balance between different stages can vary markedly in different games. This is illustrated by F.E.A.R., which makes extensive use of stencil-volume shadows. This results in a reduced pixel shading load, but heavy rasterization and depth test loads. This shows the importance of being able to reconfigure the computing resource allocation among different stages, including rasterization, which is 3.2% in two of the games but 20.1% in F.E.A.R.



**Figure 13: End-to-End Average Time Breakdowns:** shows the average time spent in each rendering stage for the three games

Figure 14 shows the time spent in each rendering stage for F.E.A.R. For the other two games, the ratios stay very similar across the tested frames, but F.E.A.R. shows significant variation. In addition, considerable variation is observed over the course of a single frame of rendering. Larrabee's cores each process an entire tile at once, then process the next and so on, leading to a reasonably uniform load over the course of the entire frame. By contrast, an immediate-mode renderer doesn't have as many ways to process pixels and primitives out of order. Further, the widely varying loads can cause different units to bottleneck at different times, unless they are over designed for the worst case.

18:10 • L. Seiler et al.



**Figure 14:** F.E.A.R. per-frame time breakdowns: this chart shows the time spent in each rendering stage for 25 widely spaced frames of F.E.A.R., which show considerable load variation.

Our conclusion is that application dependent resource balancing is not sufficient. Instead, dynamic load balancing is likely to be very important to achieving high average performance. Larrabee's entirely software scheduling algorithms provide a great deal of flexibility for adjusting load balancing algorithms.

## 6. Advanced Applications

Larrabee supports performance implementations of many other parallel applications. Section 6.1 describes how applications can be developed using traditional multi-core high level languages and tools that have been targeted to Larrabee's many-core architecture. Section 6.2 discusses Larrabee support for irregular data structures, which are common in these applications [Pharr 2006]. Sections 6.3 and 6.4 describe results of simulating rendering and other throughput applications on Larrabee.

Scalability and performance analysis in this section uses an in-house simulator that models variable configurations of Larrabee cores, threads, and memory hierarchy. This simulator is derived from proven cycle accurate simulator technology used in the design of CPU cores. Reported data is from hand coded and software threaded kernels running on this simulator.

### 6.1 Larrabee Many-Core Programming Model

The *Larrabee Native* programming model resembles the well known programming model for x86 multi-core architectures. Central to Larrabee Native programming is a complete C/C++ compiler that statically compiles programs to the Larrabee x86 instruction set. Many C/C++ applications can be recompiled for Larrabee and will execute correctly with no modification. Such application portability alone can be an enormous productivity gain for developers, especially for large legacy x86 code bases like those found in high-performance computing and numeric-intensive computing environments. Two current limitations are that application system call porting is not supported and the current driver architecture requires application recompilation.

We now discuss three important aspects of application programming for Larrabee Native: software threading, SIMD vectorization, and communication between the host and Larrabee.

Larrabee Native presents a flexible software threading capability. The architecture level threading capability is exposed as the well known POSIX Threads API (P-threads) [IEEE 2004]. We have extended the API to also allow developers to specify thread affinity with a particular HW thread or core.

Although P-threads is a powerful thread programming API, its thread creation and thread switching costs may be too high for some application threading. To amortize such costs, Larrabee Native provides a task scheduling API based on a light weight distributed task stealing scheduler [Blumofe et al. 1996]. A production implementation of such a task programming API can be found in Intel Thread Building Blocks [Reinders 2007]. Finally, Larrabee Native provides additional thread programming support via OpenMP [Chandra et al. 2000] pragmas in Larrabee Native's C/C++ compiler.

All Larrabee SIMD vector units are fully programmable by Larrabee Native application programmers. Larrabee Native's C/C++ compiler includes a Larrabee version of Intel's auto-vectorization compiler technology. Developers who need to program Larrabee vector units directly may do so with C++ vector intrinsics or inline Larrabee assembly code.

In a CPU based platform that includes a Larrabee based add-in card, Larrabee will be managed by an OS driver for that platform. In such a platform, Larrabee Native binaries are tightly paired with a host binary. Larrabee libraries provide fast message/data passing protocol to manage all memory transfers and communications between the binaries. The API supports both synchronous and asynchronous data transfers. Additionally, execution of some C/C++ standard library functions called from Larrabee application binaries must be shared with the host operating system. Specifically file I/O functions such as read/write/open/close, etc., are proxied from the Larrabee application binary back to a service that executes such functions remotely on the host OS.

Besides high throughput application programming, we anticipate that developers will also use Larrabee Native to implement higher level programming models that may automate some aspects of parallel programming or provide domain focus. Examples include C++ style programming models [Ghouloum et al. 2007], high level library APIs such as Intel® Math Kernel Library (Intel® MKL) [Chuvelev et al. 2007], and physics APIs. Existing GPGPU programming models can also be re-implemented via Larrabee Native if so desired [Buck et al. 2004; Nickolls et al. 2008].

### 6.2 Irregular Data Structure Support

Larrabee provides excellent support for high throughput applications that use irregular data structures such as complex pointer trees, spatial data structures, or large sparse n-dimensional matrices. They are supported by Larrabee's programming model, memory hierarchy, and VPU instructions.

For Larrabee applications, the multithreaded C++ code to populate, transform, or traverse these data structures follows the familiar programming methodology used on multi-core CPUs. C++ pointers, inheritance, and classes may be used to implement graph nodes. The individual nodes may have significantly different operation execution costs or code branch behavior. Because thread or task scheduling is under programmer control, tasks that operate on these data structures can be dynamically re-bundled to maintain SIMD efficiency. For example, a ray tracer's secondary reflection rays may be re-bundled differently than the primary camera rays that generated them. Finally, data structure techniques such as pre-allocated memory pools can be used to asymmetrically provide only the memory required for a given data structure node. For example, an order-independent transparency implementation may dynamically associate memory based on the number of layers per pixel, rather than pre-allocating a wasteful overestimation of the number of layers per pixel as K-buffer techniques often do [Callahan et al. 2005; Bavoil et al. 2007].



Unlike stream based architectures [Pham et al. 2005; Khailany et al. 2002], Larrabee allows but does not require direct software management to load data into different levels of the memory hierarchy. Software simply reads or writes data addresses, and hardware transparently loads data across the hierarchy. Software complexity is significantly reduced and data structures can employ hard to predict unstructured memory accesses.

In recent Nvidia GPUs, local shared memory support is provided through small (16KB on Nvidia GeForce 8) Per-Block Shared Memories (PBSMs). Each PBSM is shared by 8 scalar processors running up to 768 program instances (which Nvidia calls *threads*) within a SIMD multi-processor [Nickolls et al. 2008]. For high speed local sharing, programmers must explicitly load shared data structures into a PBSM. These are not directly shareable by instances in a different SIMD group. Similarly, order and consistency protection requires software to issue a barrier sync that is visible only within a SIMD group. To facilitate broader sharing across SIMD groups, data must be explicitly written out to higher latency GDDR memory. In contrast, all memory on Larrabee is shared by all processor cores. For Larrabee programmers, local data structure sharing is transparently supported by the coherent cached memory hierarchy regardless of the thread's processor. Protection can be provided by conventional software locks, semaphores, or critical sections.

An important aspect of handling irregular data structures is efficient scatter-gather support, so that the SIMD VPU can work on non-contiguous data. As described in Section 3, Larrabee implements VPU scatter-gather instructions which load a VPU vector register from sixteen non-contiguous memory locations. The non-contiguous data elements can reside anywhere in the large on-die cache, without suffering memory access penalty. This significantly reduces programmer data management overhead. We have observed an average of almost 3x performance gain in Larrabee from hardware support of scatter-gather, compared to software scatter-gather, for basic sparse matrix compute kernels, such as sparse matrix-vector multiply. Algorithms requiring irregular data structures also benefit from Larrabee instructions such as count bits, bit scan, and packed loads and stores.

### 6.3 Extended Rendering Applications

The Larrabee graphics rendering pipeline is itself a Larrabee Native application. Because it is software written with high level languages and tools, it can easily be extended to add innovative rendering capabilities. Here we briefly discuss three example extensions of the graphics pipeline that we are studying. Future implementations could evolve towards a fully programmable graphics pipeline as outlined by Pharr [2006].

**Render Target Read:** Because Larrabee's graphics rendering pipeline employs a software frame buffer, we can enable additional programmer access to those data structures. More specifically, a trivial extension to the Larrabee rendering pipeline would be to allow pixel shaders to directly read previously stored values in render targets. Such a capability could serve a variety of rendering applications, including programmer defined blending operations, single-pass tone mapping, and related functions.


**Order Independent Transparency:** Presently 3D application developers must either depth sort translucent models in their application every frame prior to rendering or else implement multi-pass algorithms such as depth peeling [Wexler et al. 2005] to achieve correct inter-model transparency. Neither method allows the kinds of post-rendering area effects that are possible

with opaque models. Figure 15 illustrates artifacts that occur if such effects are applied after merging the translucent surfaces.



*Figure 15: Transparency without and with pre-resolve effects: the left image sorts the geometry and resolves before applying a fog patch. The right image applies the fog patch to the translucent surfaces and then resolves. The fog is visible through the wing in the right image, but not in the left image (Dragon models designed and created by Jeffery A. Williams and Glen Lewis.)*

Larrabee can support order independent transparency (OIT) with no additional dedicated logic by storing multiple translucent surfaces in a per-pixel spatial data structure. After rendering the geometry, we can perform effects on the translucent surfaces, since each surface retains its own depth and color, before sorting and resolving the fragment samples stored per pixel.



*Figure 16: Irregular Z-Buffer sample frame: this method uses an irregular spatial data structure to produce alias-free shadowing. Like the transparency example in figure 15, the data structure is tightly integrated with the rendering pipeline. The renderer constructs the shadowmap and then treats it as a special class of frame buffer. (Skeleton model by TurboSquid.)*

**Irregular Shadow Mapping:** Shadow mapping is a popular real-time shadow approximation technique, but most implementations are plagued by visually displeasing aliasing artifacts. A variety of heuristics have been proposed in an attempt to reduce artifacts [Akenine-Möller & Haines 2002; Bookout 2007]. Irregular shadow mapping (ISM) offers an exact solution to this problem and places no additional burden on the application programmer [Aila & Laine 2004; Johnson et al. 2005].

To implement ISM, we dynamically construct a spatial data structure in the light view using depth samples captured in the camera view. We again customize Larrabee's all software graphics pipeline by adding a stage that performs light view ISM rasterization against ISM's spatial data structure. Because the shadow map is computed at exact positions, the resulting shadow map is alias free. This technique can be used to achieve real-time hard shadowing effects, as shown in Figure 16, or as the foundation for real-time soft shadowing effects.

6.4 Other Throughput Computing Applications

Larrabee is also suitable for a wide variety of non-rasterization based throughput applications. The following is a brief discussion of the observed scalability and characteristics of several examples.




Figure 17: Game Physics Scalability Performance: this shows that the Larrabee architecture is scalable to meet the growing performance needs of interactive rigid body, fluid, and cloth simulation algorithms and some commonly used collision kernels.

**Game Physics:** We have performed detailed scalability simulation analysis of several game physics workloads on various configurations of Larrabee cores. Figure 17 shows scalability of some widely used game physics benchmarks and algorithms for rigid body, fluid, and cloth. We achieve better than 50% resource utilization using up to 64 Larrabee cores, and achieve near-linear parallel speedup in some cases. The game rigid body simulation is based on the popular “castle” destruction scene with 10K objects. Scalability plots for Sweep-and-Prune [Cohen et al. 1995] and GJK [Gilbert et al. 1988] distance algorithms are included since they are some of the most commonly used collision detection routines. Game fluid simulation is based on the smoothed particle hydrodynamics (SPH) algorithm. We used a mass spring model and Verlet integration for our game cloth simulation [Jacobsen 2001]. Bader et al. [2008] provide details on the implementation and scalability analysis for these game physics workloads




Figure 18: Real time ray tracing on Larrabee: cropped from a 1Kx1K sample image that requires ~4M rays. The ray tracer was implemented in C++ with some hand-coded assembly code for key routines like ray intersection. Kd-trees are typically 25MB and are built dynamically per frame. Primary and reflection rays are tested in 16 ray bundles. Nearly all 234K triangles are visible to primary or reflection rays. (Bar Carta Blanca model by Guillermo M Leal Laguno, courtesy of Cornell University.)

**Real Time Ray Tracing:** The highly irregular nature of spatial data structures used in Whitted-style real-time ray tracers benefit from Larrabee’s general purpose memory hierarchy, relatively short pipeline, and VPU instruction set. Here we used SIMD 16 packet ray tracing traversing through a kd-tree. For the complete workload, we observe that a single Intel Core 2 Duo processor requires 4.67x more clock cycles than a single Larrabee core, which shows the effectiveness of the Larrabee instruction set and wide SIMD. Results are even better for small kernels. For example, the intersection test of 16 rays to 1 triangle takes 47 cycles on a single Larrabee core. The same test takes 257 Core 2 Duo processor cycles. Figure 18 shows a 1024x1024 frame of the bar scene with 234K triangles, 1 light source, 1 reflection level, and typically 4M rays per frame. Figure 19 compares performance for Larrabee with an instance of the ray tracer running on an Intel Xeon® processor 2.6GHz with 8 cores total. Shevtsov et al. [2007] and Reshetov et al. [2005] describe details of this implementation.





Figure 19: Real time ray tracing scalability: this graph compares different numbers of Larrabee cores with a nominal 1GHz clock speed to an Intel Xeon processor 2.6GHz with 8 cores total. The latter uses 4.6x more clock cycles than are required by 8 Larrabee cores due to Larrabee’s wide VPU and vector instruction set. Figure 18 describes the workload for these tests.

**Image and Video Processing:** The Larrabee architecture is suitable for many traditional 2D image and video analysis applications. Native implementations of traditional 2D filtering functions (both linear and non-linear) as well as more advanced functions, like video cast indexing, sports video analysis, human body tracking, and foreground estimation offer significant scalability as shown in Figure 20. Biomedical imaging represents an important subset of this processing type. Medical imaging needs such as back-projection, volume rendering, automated segmentation, and robust deformable registration, are related yet different from those of consumer imaging and graphics. Figure 20 also includes scalability analysis of iso-surface extraction on a 3D volume dataset using the marching cubes algorithm.

**Physical Simulation:** Physical simulation applications use numerical simulation to model complex natural phenomena in movies and games, such as fire effects, waterfalls in virtual worlds, and collisions between rigid or deformable objects. Large data-sets, unstructured control-flow and data accesses often make these applications more challenging to scale than traditional streaming applications. Looking beyond interactive game physics, we also analyzed applicability of Larrabee architecture for the broader class of entertainment physics including offline movie-industry effects and distributed real-time virtual-world simulation. Specific simulation results based on Stanford’s PhysBAM are shown in Figure 20 and illustrate very good scalability for production fluid, production cloth, and production face. Implementation and scalability analysis details are described by Hughes et al. [2007].



**Figure 20:** Scalability of select non-graphics applications and kernels: Larrabee's general-purpose many-core architecture delivers performance scalability for various non-graphics visual and throughput computing workloads and common HPC kernels.

Larrabee is also highly scalable for non-visual throughput applications, as shown in Figure 20. Larrabee's highly-threaded x86 architecture benefits traditional enterprise throughput computing applications, such as *text indexing*. Its threading, together with its wide-SIMD IEEE-compliant double-precision support, makes it well positioned for financial analytics, such as *portfolio management*. Internal research projects have proven Larrabee architecture scalability for many traditional high performance computing (HPC) workloads and well-known HPC kernels such as *3D-FFT* and *BLAS3* (with dataset larger than on-die cache). More details are described by Chen et al. [2008].

## 7. Conclusions

We have described the Larrabee architecture, which uses multiple x86-based CPU cores, together with wide vector processor units and some fixed function logic, to achieve high performance and flexibility for interactive graphics and other applications. We have also described a software renderer for the Larrabee architecture and a variety of other throughput applications, with performance and scalability analysis for each. Larrabee is more programmable than current GPUs, with fewer fixed function units, so we believe that Larrabee is an appropriate platform for the convergence of GPU and CPU applications.

We believe that this architecture opens a rich set of opportunities for both graphics rendering and throughput computing. We have observed a great deal of convergence towards a common core of computing primitives across the workloads that we analyzed on Larrabee. This underlying workload convergence [Chen et al. 2008] implies potential for a common programming model, a common run-time, and a native Larrabee implementation of common compute kernels, functions, and data structures.

**Acknowledgements:** The Larrabee project was started by Doug Carmean and Eric Sprangle, with assistance from many others, both inside and outside Intel. The authors wish to thank many people whose hard work made this project possible, as well as many who helped with this paper. Workload implementation and data analysis were provided by Jeff Boody, Dave Bookout, Jatin Chhugani, Chris Gorman, Greg Johnson, Danny Lynch, Oliver Macquelin, Teresa Morrison, Misha Smelyanskiy, Alexei Soupikov, and others from Intel's Application Research Lab, Software Systems Group, and Visual Computing Group.

## References

- AKENINE-MÖLLER, T., HAINES, E. 2002. *Real-Time Rendering*. 2nd Edition. A. K. Peters.
- AILA, T., LAINE, S. 2004. Alias-Free Shadow Maps. In *Proceedings of Eurographics Symposium on Rendering 2004*, Eurographics Association, 161-166.
- ALPERT, D., AVNON, D. 1993. Architecture of the Pentium Microprocessor. *IEEE Micro*, v. 13, n. 3, 11-21. May 1993.
- AMD. 2007. Product description web site: [ati.amd.com/products/Radeonhd3800/specs.html](http://ati.amd.com/products/Radeonhd3800/specs.html).
- BADER, A., CHHUGANI, J., DUBEY, P., JENKINS, S., MORRISON T., RAGOZIN, D., SMELYANSKIY. 2008. Game Physics Performance On Larrabee Architecture. Intel whitepaper, available in August, 2008. Web site: [techresearch.intel.com](http://techresearch.intel.com).
- BAVOIL, L., CALLAHAN, S., LEFOTIN, A., COMBA, J., SILVA, C. 2007. Multi-fragment effects on the GPU using the k-buffer. In *Proceedings of the 2007 Symposium on Interactive 3D Graphics and Games* (Seattle, Washington, April 30 - May 02, 2007). I3D 2007. ACM, New York, NY, 97-104.
- BLUMOF, R., JOERG, C., KUSZMAUL, B., LEISERSON, C., RANDALL, K., ZHOU, Y. Aug. 25, 1996. Cilk: An Efficient Multithreaded Runtime System. *Journal of Parallel and Distributed Computing*, v. 37, i. 1, 55-69.
- BLYTHE, D. 2006. The Direct3D 10 System. *ACM Transactions on Graphics*, 25, 3, 724-734.
- BOOKOUT, D. July, 2007. Shadow Map Aliasing. Web site: [www.gamedev.net/reference/articles/article2376.asp](http://www.gamedev.net/reference/articles/article2376.asp).
- BUCK, I., FOLEY, T., HORN, D., SUGERMAN, J., FATAHALIAN, K., HOUSTON, M., AND HANRAHAN, P. 2004. Brook for GPUs: stream computing on graphics hardware. *ACM Transactions on Graphics*, v. 23, n. 3, 777-786.
- CALLAHAN, S., IKES, M., COMBA, J., SILVA, C. 2005. Hardware-assisted visibility sorting for unstructured volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 11, 3, 285-295.
- CHANDRA, R., MENON, R., DAGUM, L., KOHR, D., MAYDAN, D., McDONALD, J. 2000. *Parallel Programming in OpenMP*. Morgan Kaufman.
- CHEN, M., STOLL, G., IGEHY, H., PROUDFOOT, K., HANRAHAN P. 1998. Simple models of the impact of overlap in bucket rendering. In *Proceedings of the ACM SIGGRAPH / EUROGRAPHICS Workshop on Graphics Hardware* (Lisbon, Portugal, August 31 - September 01, 1998). S. N. Spencer, Ed. HWS '98. ACM, New York, NY, 105-112.
- CHEN, Y., CHHUGANI, J., DUBEY, P., HUGHES, C., KIM, D., KUMAR, S., LEE, V., NGUYEN A., SMELYANSKIY, M. 2008. Convergence of Recognition, Mining, and Synthesis Workloads and its Implications. In *Proceedings of HETI*, v. 96, n. 5, 790-807.
- CHUVELEV, M., GREER, B., HENRY, G., KUZNETSOV, S., BURYLOV, I., SABANIN, B. Nov. 2007. Intel Performance Libraries: Multi-core ready Software for Numeric Intensive Computation. *Intel Technology Journal*, v. 11, i. 4, 1-10.
- COHEN, J., LIN, M., MANOCHA, D., PONAMGI, D. 1995. I-COLLIDE: An Interactive and Exact Collision Detection System for Large-Scale Environments. In *Proceedings of 1995 Symposium on Interactive 3D Graphics*, SI3D '95. ACM, New York, NY, 189-196.

18:14 • L. Seiler et al.

- ELDRIDGE, M. 2001. *Designing Graphics Architectures Around Scalability and Communication*. PhD thesis, Stanford.
- FOLEY, J., VAN DAM, A., FLEINER, S., HUGHES, J. 1996. *Computer Graphics: Principles and Practice*. Addison Wesley.
- FUCHS, H., POULTON, J., EYLES, J., GREER, T., GOLDFEATHER, J., ELLSWORTH, D., MOLNAR, S., TURK, G., TEIBBS, B., ISRAEL, I. 1989. Pixel-planes 5: a heterogeneous multiprocessor graphics system using processor-enhanced memories. In *Computer Graphics (Proceedings of ACM SIGGRAPH 89)*, ACM, 79-88.
- GHULOUM, A., SMITH, T., WU, G., ZHOU, X., FANG, J., GUO, P., SO, B., RAJAGOPALAN, M., CHEN, Y., CHEN, B. November 2007. Future-Proof Data Parallel Algorithms and Software on Intel Multi-Core Architectures. *Intel Technology Journal*, v. 11, i. 04, 333-348.
- GILBERT, E., JOHNSON, D., KEERTHI, S. 1988. A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal of Robotics and Automation*, 4, 2, 193-203.
- GPGPU. 2007. GPGPU web site: [www.gpgpu.org](http://www.gpgpu.org).
- GREENE, N. 1996. Hierarchical polygon tiling with coverage masks. In *Proceedings of ACM SIGGRAPH 93*, ACM Press/ACM SIGGRAPH, New York, NY, Computer Graphics Proceedings, Annual Conference Series, ACM, 65-64.
- GROCHOWSKI, E., RONEN, R., SHEN, J., WANG, H. 2004. Best of Both Latency and Throughput. *2004 IEEE International Conference on Computer Design (ICCD '04)*, 236-243.
- GWENNAP, L. 1995. Intel's P6 Uses Decoupled Superscalar Design. *Microprocessor Report*, v. 9, n. 2, Feb. 16, 1995.
- HSEIH, E., PENTKOWSKI, V., PIAZZA, T. 2001. ZR: A 3D API Transparent Technology For Chunk Rendering. In *Proceedings of the 34th Annual ACM/IEEE International Symposium on Microarchitecture* (Austin, Texas, December 01 - 05, 2001). International Symposium on Microarchitecture. IEEE Computer Society, Washington, DC, 284-291.
- HUGHES, C. J., GRZESZCZUK, R., SIFAKIS, E., KIM, D., KUMAR, S., SELLE, A. P., CHITUGANI, J., HOLLIMAN, M., AND CHEN, Y. 2007. Physical simulation for animation and visual effects: parallelization and characterization for chip multiprocessors. In *Proceedings of the 34th Annual international Symposium on Computer Architecture* (San Diego, California, USA, June 09 - 13, 2007). ISCA '07. ACM, New York, NY, 220-231.
- IEEE Std. 1003.1, 2004 Edition. Standard for Information Technology – Portable Operating System Interface (POSIX) System Interfaces. *The Open Group Technical Standard Base Specifications*. Issue 6.
- JACOBSEN, T. 2001. Advanced Character Physics. *Proc. Game Developers Conference 2001*, 1-10.
- JOHNSON, G.S., LEE, J., BURNS, C.A., MARK, W.R. 2005. The irregular Z-buffer: Hardware acceleration for irregular data structures. *ACM Transactions on Graphics*, 24, 4, 1462-1482.
- KELLEY, M., GOULD, K., PEASE, B., WINNER, S., YEN, A. 1994. Hardware accelerated rendering of CSG and transparency. In *Proceedings of SIGGRAPH 1994*, ACM Press/ACM SIGGRAPH, New York, NY, Computer Graphics Proceedings, Annual Conference Series, ACM, 177-184.
- KELLEY, M., WINNER, S., GOULD, K. 1992. A Scalable Hardware Render Accelerator using a Modified Scanline Algorithm. In *Computer Graphics (Proceedings of ACM SIGGRAPH 1992)*, SIGGRAPH '92. ACM, New York, NY, 241-248.
- KESSENICH, J., BALDWIN, D., ROST, R. The OpenGL Shading Language. Version 1.1. Sept. 7, 2006. Web site: [www.opengl.org/registry/doc/GLSLangSpec.Full.1.20.8.pdf](http://www.opengl.org/registry/doc/GLSLangSpec.Full.1.20.8.pdf)
- KHAILANY, B., DALLY, W., RIXNER, S., KAPASI, U., MATTISON, P., NAMKOONG, J., OWENS, J., TOWLES, B., CHANG, A. 2001. Imagine: Media Processing with Streams. *IEEE Micro*, 21, 2, 35-46.
- KONGETIRA, P., ANGARAN, K., OLUKOTUN, K. Mar/Apr 2005. Niagara: A 32-way multithreaded SPARC Processor. *IEEE Micro*, v. 25, i. 2, 21-29.
- LAKE, A. 2005. Intel Graphics Media Accelerator Series 900 Developer's Guide. Version 2.0. Web site: [download.intel.com/ids/gma/Intel\\_915G\\_SDG\\_Feb05.pdf](http://download.intel.com/ids/gma/Intel_915G_SDG_Feb05.pdf).
- LLOYD, B., GOVINDARAJU, N., MOLNAR, S., MANOCHIA, D. 2007. Practical logarithmic rasterization for low-error shadow maps. In *Proceedings of the 22nd ACM SIGGRAPH/EUROGRAPHICS Symposium on Graphics Hardware*, 17-24.
- MARK, W., GLANVILLE, S., AKELEY, K., KILGARD, M. 2003. Cg: A System for Programming Graphics Hardware in a C-like Language. *ACM Transactions on Graphics*, v. 22, n. 3, 896-907.
- MICROSOFT. 2007. Microsoft Reference for HLSL. Web site: [msdn2.microsoft.com/en-us/library/bb509638.aspx](http://msdn2.microsoft.com/en-us/library/bb509638.aspx).
- MOLNAR, S., COX, M., ELLSWORTH, D., AND FUCHS, H. 1994. A Sorting Classification of Parallel Rendering. *IEEE Computer Graphics and Applications*, v.14, n. 4, July 1994, 23-32.
- MOLNAR, S., EYLES, J., POULTON, J. 1992. Pixelflow: High Speed Rendering Using Image Composition. *Computer Graphics (Proceedings of SIGGRAPH 92)*, v. 26, n. 2, 231-240.
- MOREIN, S. 2000. ATI Radeon HyperZ Technology. Presented at *Graphics Hardware 2000*. Web site: [www.graphicshardware.org/previous/www\\_2000/presentations/ATHot3D.pdf](http://www.graphicshardware.org/previous/www_2000/presentations/ATHot3D.pdf).
- NICKOLLS, J., BUCK, I., GARLAND, M. 2008. Scalable Parallel Programming with CUDA. *ACM Queue*, 6, 2, 40-53.
- NVIDIA. 2008. Product description web site: [www.nvidia.com/object/geforce\\_family.html](http://www.nvidia.com/object/geforce_family.html).
- OWENS, J., LUEBKE, D., GOVINDARAJU, N., HARRIS, M., KRUGER, J., LERDOIN, A., PURCELL, T. 2007. A Survey of General Purpose Computation on Graphics Hardware. *Computer Graphics Forum*, v.26, n. 1, 80-113.
- PIAM D., ASANO, S., BOLLIGER, M., DAY, M., HOFSTEE, H., JOHNS, C., KAHLE, J., KAMEYAMA, A., KEATY, J., MASUBUCHI, Y., RILEY, M., SHIPPY, D., STASIASK, D., SUZUODI, M., WANG, M., WARNOCK, J., WREITZEL, S., WENDEL, D., YAMAZAKI, T., YAZAWA, K. 2005. The Design and Implementation of a First Generation CELL Processor. *IEEE International Solid-State Circuits Conference*, 184-186.
- PHARR, M. 2006. Interactive Rendering in the Post-GPU Era. Presented at *Graphics Hardware 2006*. Web site: [www.pharr.org/matt/](http://www.pharr.org/matt/).
- PINEDA, J. 1988. A Parallel Algorithm for Polygon Rasterization. In *Computer Graphics (Proceedings of ACM SIGGRAPH 88)*, 22, 4, 17-20.
- POWERVR. 2008. Web site: [www.imgtec.com/powervr/products/Graphics/index.asp](http://www.imgtec.com/powervr/products/Graphics/index.asp).
- POLLACK, F. 1999. New Microarchitecture Challenges for the Coming Generations of CMOS Process Technologies. *Micro32*.

ACM Transactions on Graphics, Vol. 27, No. 3, Article 18, Publication date: August 2008.



- REINDERS, J., 2007. *Intel Threading Building Blocks: Outfitting C++ for Multi-core Processor Parallelism*. O'Reilly Media, Inc.
- RESHETOV A., SOUPIKOV, A., HURLEY, J. 2005. Multi-level Ray Tracing Algorithm. *ACM Transactions on Graphics*, 24, 3, 1176-1185.
- ROST, R. 2004. *The OpenGL Shading Language*. Addison Wesley.
- SILVETSOV, M., SOUPIKOV, A., KAPUSTIN, A. 2007. Ray-Triangle Intersection Algorithm for Modern CPU Architectures. In *Proceedings of GraphiCon 2007*, 33-39.
- STEVENS, A. 2006. ARM Mali 3D Graphics System Solution. Web site: [www.arm.com/miscPDFs/16514.pdf](http://www.arm.com/miscPDFs/16514.pdf).
- STOLL, G., ELDRIDGE, M., PATTERSON, D., WEBB, A., BERMAN, S., LEVY, R., CAYWOOD, C., TAVEIRA, M., HUNT, S., HANRAHAN, P. 2001. Lightning 2: A High Performance Display Subsystem for PC Clusters. In *Computer Graphics (Proceedings of ACM SIGGRAPH 2001)*, ACM, 141-148.
- TORBORC, J., KAJIYA, J. 1996. Talisman Commodity Realtime 3D Graphics for the PC. In *Proceedings of ACM SIGGRAPH 1996*, ACM Press/ACM SIGGRAPH, New York. Computer Graphics Proceedings, Annual Conference Series, ACM, 353-363.
- WEXLER, D., GRITZ, L., ENDERTON, E., RICE, J. 2005. GPU-accelerated high-quality hidden surface removal. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on Graphics Hardware* (Los Angeles, California, July 30 - 31, 2005). HWWS '05, ACM, New York, NY, 7-14.





# EXHIBIT 4

**From:** Greg Brandeau  
**To:** Gelsinger, Patrick P  
**CC:** Cooper, Don  
**Sent:** 10/10/2008 10:46:20 PM  
**Subject:** Re: Intel recruiting Pixar employees

Hi Pat,

Thanks for taking care of this. I owe Don a return call which I didn't get to today but I will on Monday.

Regards,

-gtb

Gelsinger, Patrick P wrote:

> Greg -

>

> I apologize this has happened. As you know, we greatly value your partnership.

>

> As you can understand, we're very aggressively driving our graphics/media recruiting. We've used a variety of direct and indirect recruiters to expedite the build-up of our resources in this area.

>

> Based on your inquiry, we've immediately directed our recruiters that Intel will not proactively pursue any Pixar employees going forward. I've copied my director of recruiting Don Cooper on this message and he is responsible for implementing this policy. Feel free to have anyone at Pixar interact with him directly for any specific situations or concerns. Don will also be investigating the past situations and provide me a bit more background as well.

>

> Regards,

> Pat.

>

> -----Original Message-----

> From: Greg Brandeau [mailto:brandeau@pixar.com]

> Sent: Thursday, October 09, 2008 5:31 PM

> To: Gelsinger, Patrick P

> Cc: Prajapati, Ranna H; Christopher Ford

> Subject: Intel recruiting Pixar employees

>

> Hi Pat,

>

> Over the past 6-9 months, Intel has made offers to 3 engineers who work in our RenderMan team. Given that Pixar and Intel do have a close relationship, could you ask the recruiters inside of Intel to please stop recruiting our people? I am not saying the crazy thing that people at Pixar shouldn't be allowed to apply to Intel job postings, I am simply asking that you don't cold call our people. Especially Dana's team which is only 15 people in the first place!

>

> Thanks!

>

> -gtb

>

40006DOC000516  
CONFIDENTIAL

# EXHIBIT 5

**From:** Prajapati, Ranna H  
**To:** Connelly, Jay; Gelsinger, Patrick P; Cooper, Don  
**CC:** Garbus, Elliot D; Georgiopoulos, Christos; Mark, William R  
**Sent:** 10/14/2008 9:04:45 AM  
**Subject:** RE: Update on Pixar discussions...

Hi Jay,

Thank you for very much for speaking with Greg and addressing his hiring concerns below... much appreciated.

As for the perception of moving Ram from Pixar to DWA, I've been working on that one for the last few months and have spoken to Ram and Pixar's Renderman team several times to assure them there is no breach of confidentiality regarding Pixar's rendering technology. Please let me know what additional assistance you need from me on this front...

Regards,

**Ranna Prajapati**

Enterprise Account Manager

Enterprise Solution Sales

Office: 916.356.2758

Mobile: 916.769.2231

E-mail: [ranna.h.prajapati@intel.com](mailto:ranna.h.prajapati@intel.com)

**From:** Connelly, Jay  
**Sent:** Tuesday, October 14, 2008 8:30 AM  
**To:** Gelsinger, Patrick P; Cooper, Don; Prajapati, Ranna H  
**Cc:** Garbus, Elliot D; Georgiopoulos, Christos; Mark, William R  
**Subject:** Update on Pixar discussions...

Pat, Don and Ranna –

I wanted to give you an update on a call I received yesterday from Greg Brandeau at Pixar.

On the hiring front, I think we are fine. I explained that we were not actively targeting Pixar employees and that the 2 rock stars we hired in AVC were based on long standing personal relationships with Matt Pharr and Craig Kolb. I think he understood (or seemed to).

The incident that triggered the mail to Pat was a 3<sup>rd</sup> employee (from the same group) that Bill Mark was talking to. I knew nothing about this but assured Greg that I'd talk with Bill and let him know of the sensitivities. I have, Bill knows and its all good.

The last remaining issue here has to do with Greg's perception that we moved an AE from Pixar to DWA and I'll work with Christos to resolve. Greg's concern was that Ram learned everything he knows about rendering from Pixar and feels violated that we would use that knowledge to help a competitor.

Cheers –

J

40006DOC000523  
 CONFIDENTIAL

# EXHIBIT 6

**From:** Greg Brandeau  
**To:** Gelsinger, Patrick P  
**CC:** Cooper, Don; Connelly, Jay; Pattani, Paresh G; Prajapati, Ranna H; Dana Batali; Christopher Ford  
**Sent:** 10/16/2008 11:49:54 PM  
**Subject:** Re: Intel recruiting Pixar employees + Dreamworks

Hi Pat and Don,

Here is an example of recruiting that happened today.

This highlights one of my other topics, potential IP leakage from Pixar, via Intel, to Dreamworks. I have a call scheduled with Paresh in the morning about that topic.

-gtb

>  
 > ----- Forwarded Message -----  
 > From: "Jeanine Hughes" <jeanine.hughes@intel.com>  
 > To: "Chris Ford" <cford81@comcast.net>  
 > Sent: Thursday, October 16, 2008 3:03:43 PM GMT -08:00 US/Canada Pacific  
 > Subject: Animation Opportunities at Intel  
 >  
 > LinkedIn  
 > INMAIL: YOU HAVE A NEW MESSAGE  
 > From: Jeanine Hughes  
 > Date: 10/16/2008  
 > Subject: Animation Opportunities at Intel  
 > Hi Chris We are currently seeking bright, talented software  
 > graphics/animation engineers to work on our project with DreamWorks.  
 > Below is an overview for one of our new positions. Would you know  
 > anyone that may be interested in learning more about this opportunity?  
 > Jeanine Hughes Strategic Recruiting Team Intel Corporation  
 > jeanine.hughes@intel.com <http://www.linkedin.com/in/hughesjeanine>  
 > Senior Software Engineer Description: Join a team of the brightest  
 > minds in computer-generated animation, graphics acceleration  
 > technology, multicore hardware, and software engineering onsite at  
 > DreamWorks Animation. Help redefine how the world develops leading 3D  
 > graphics and visualization software. Responsibilities include  
 > optimizing existing production content-creation applications and  
 > developing new highly-parallel production systems. This position also  
 > offers an opportunity to advance the state-of-the-art by implementing  
 > advanced applications hosted on novel computing platforms.  
 > View/reply to this message  
 > Don't want to receive e-mail notifications? Adjust your message settings.  
 > © 2008, LinkedIn Corporation

40006DOC000527  
CONFIDENTIAL

# EXHIBIT 7



**From:** Pattani, Paresh G  
**To:** Pattani, Paresh G; Cooper, Don; Prajapati, Ranna H; Sathis, Paul; Connelly, Jay  
**CC:** Dickenson, Christina L  
**Sent:** 10/17/2008 11:42:21 AM  
**Subject:** My notes from meeting with Greg Brandeau

Folks,

Following are the notes from my meeting with Greg Brandeau, SVP Technology, Pixar:

Two topics:

**Intel proactively recruiting Pixar employees:** Greg was upset that after talking with Pat Gelsinger and reaching an agreement with Intel not to proactively recruit Pixar employees we are still continue to do it. I assured him that we will provide strict guidelines to our staffing recruiters that they should not proactively go after Pixar employees. Greg understands that Intel cannot help it if Pixar employees approach us by themselves.

**IP Transfer through Ram Ramanujam (SSG employee previously worked on Pixar account):** I informed him that last time Ram had seen source code was 4 years back on-site at Pixar. Also informed him that Ram is working on Animation not on Rendering and focusing on Intel strengths in parallelization. He feels that Ram may have learned something about how they use memory to improve scaling. He was not concerned about source code IP migrating to DWA but more concerned about algorithmic learning's. He is concerned about **future work with Intel on Larrabee visualization**. He wants to work with SSG but concerned about same thing might happen again. I suggested that they should put specific AE names in the contract and bar them from working on specified competitor's products for specified time. He seemed to be fine with that but wants to talk to Renee James directly to express his opinion. **He will directly try to contact Renee.**

Let me if you have any questions.

Paresh

-----Original Message-----

From: Pattani, Paresh G  
 Sent: Friday, October 17, 2008 10:15 AM  
 To: Cooper, Don; Prajapati, Ranna H; Sathis, Paul; Connelly, Jay  
 Cc: Dickenson, Christina L  
 Subject: RE: Intel recruiting Pixar employees + Dreamworks

Don,

Completely understand what you mention below. I think we need to instruct all recruiters not to proactively approach Pixar employees. If Pixar employees approach us that is a different matter. This is a sensitive area and we should do everything to diffuse the situation.

Can you follow up on talking to staffing? We will also do from our end as well.

Jay, are you okay to instruct your staffing consultants?

Paresh

-----Original Message-----

From: Cooper, Don

Sent: Friday, October 17, 2008 9:22 AM

To: Prajapati, Ranna H; Sathis, Paul; Pattani, Paresh G; Connelly, Jay

40006DOC000583  
CONFIDENTIAL

Cc: Dickenson, Christina L

Subject: RE: Intel recruiting Pixar employees + Dreamworks

Hi Paresh, Looking at the LinkedIn InMail that Jeanine Hughes sent we may need to discuss this. Chris (The Pixar Employee) has on his linked in account that he is open to talk about jobs. LinkedIn is a third party social media board and anyone who wants to look for opportunities can elect in. It would be very hard to control this. Let me get Christy Dickenson's (Strategic Recruiting Manager) opinion on this. I'm sure we can ask the recruiters to be careful on this but when your doing large mailing lists it may be hard to control..  
Thanks, Don

-----Original Message-----

From: Prajapati, Ranna H

Sent: Friday, October 17, 2008 9:11 AM

To: Sathis, Paul; Pattani, Paresh G; Cooper, Don; Connelly, Jay

Subject: RE: Intel recruiting Pixar employees + Dreamworks

Hi Paresh,

Yes, thank you very much for doing so and let me know what additional support/action items you need from me moving forward so we can put this to rest once and for all!

Regards,

Ranna

-----Original Message-----

From: Sathis, Paul

Sent: Friday, October 17, 2008 9:01 AM

To: Pattani, Paresh G; Prajapati, Ranna H; Cooper, Don; Connelly, Jay

Subject: RE: Intel recruiting Pixar employees + Dreamworks

Thank you Paresh

Yes, we need a clear direct message to any subs as well.

PaulS

-----Original Message-----

From: Pattani, Paresh G

Sent: Friday, October 17, 2008 8:59 AM

To: Prajapati, Ranna H; Cooper, Don; Connelly, Jay

Cc: Sathis, Paul

Subject: RE: Intel recruiting Pixar employees + Dreamworks

I am talking with Greg shortly. We will talk about this issue as well. Staffing consultants are sort of independent and they dig around for candidates everywhere. I will let them know that they don't proactively approach Pixar employees.

Paresh

40006DOC000584  
CONFIDENTIAL

-----Original Message-----

From: Prajapati, Ranna H

Sent: Friday, October 17, 2008 8:29 AM

To: Cooper, Don; Connelly, Jay; Pattani, Paresh G; Prajapati, Ranna H

Cc: Sathis, Paul

Subject: FW: Intel recruiting Pixar employees + Dreamworks

Hi Everyone,

I thought we'd agreed not to pro-actively recruit for such type of positions at Pixar? How do we plan on addressing this?

Regards,

Ranna Prajapati

Enterprise Account Manager

Enterprise Solution Sales

Office: 916.356.2758

Mobile: 916.769.2231

E-mail: ranna.h.prajapati@intel.com

-----Original Message-----

From: Greg Brandeau [<mailto:brandeau@pixar.com>]

Sent: Thursday, October 16, 2008 11:50 PM

To: Gelsinger, Patrick P

Cc: Cooper, Don; Connelly, Jay; Pattani, Paresh G; Prajapati, Ranna H; Dana Batali; Christopher Ford

Subject: Re: Intel recruiting Pixar employees + Dreamworks

Hi Pat and Don,

Here is an example of recruiting that happened today.

This highlights one of my other topics, potential IP leakage from Pixar, via Intel, to Dreamworks. I have a call scheduled with Paresh in the morning about that topic.

-gtb

&gt;

&gt; ----- Forwarded Message -----

> From: "Jeanine Hughes" <[jeanine.hughes@intel.com](mailto:jeanine.hughes@intel.com)>> To: "Chris Ford" <[cford81@comcast.net](mailto:cford81@comcast.net)>

&gt; Sent: Thursday, October 16, 2008 3:03:43 PM GMT -08:00 US/Canada Pacific

40006DOC000585  
CONFIDENTIAL

> Subject: Animation Opportunities at Intel

>

> LinkedIn

> INMAIL: YOU HAVE A NEW MESSAGE

> From: Jeanine Hughes

> Date: 10/16/2008

> Subject: Animation Opportunities at Intel

> Hi Chris We are currently seeking bright, talented software

> graphics/animation engineers to work on our project with DreamWorks.

> Below is an overview for one of our new positions. Would you know

> anyone that may be interested in learning more about this opportunity?

> Jeanine Hughes Strategic Recruiting Team Intel Corporation

> jeanine.hughes@intel.com <http://www.linkedin.com/in/hughesjeanine>

> Senior Software Engineer Description: Join a team of the brightest

> minds in computer-generated animation, graphics acceleration

> technology, multicore hardware, and software engineering onsite at

> DreamWorks Animation. Help redefine how the world develops leading 3D

> graphics and visualization software. Responsibilities include

> optimizing existing production content-creation applications and

> developing new highly parallel production systems. This position also

> offers an opportunity to advance the state-of-the-art by implementing

> advanced applications hosted on novel computing platforms.

> View/reply to this message

> Don't want to receive e-mail notifications? Adjust your message settings.

> © 2008, LinkedIn Corporation

40006DOC000586  
CONFIDENTIAL

# EXHIBIT 8

**From:** Hurley, Jim <jim.hurley@intel.com>  
**Sent:** Wednesday, November 5, 2008 7:43 AM (GMT)  
**To:** Prajapati, Ranna H <ranna.h.prajapati@intel.com>  
**Subject:** RE: Status at Pixar

---

I'm in Europe, meeting on the 'phone will be problematic. Jim.

**From:** Prajapati, Ranna H  
**Sent:** Tuesday, November 04, 2008 3:54 PM  
**To:** Hurley, Jim  
**Subject:** RE: Status at Pixar

Hi Jim.... Update meaning? I can give you a phone call tomorrow to discuss.... Basically Renee didn't really help, but then again, it didn't get worse, so I don't know WHAT to think of this...

-Ranna

---

**From:** Hurley, Jim  
**Sent:** Tuesday, November 04, 2008 10:30 AM  
**To:** Prajapati, Ranna H  
**Subject:** Re: Status at Pixar

Hi Ranna, any update on this, our discussions at ILM basically have us going full circle back to Pixar. Jim.

---

**From:** Prajapati, Ranna H  
**To:** Rattner, Justin; Hurley, Jim  
**Cc:** Schutz, Joseph  
**Sent:** Tue Oct 28 14:56:19 2008  
**Subject:** RE: Status at Pixar

Hi Justin... as a courtesy call, is it possible for you to touch base with him as well? Just a thought as I believe Greg would appreciate this as well....

Thanks,  
Ranna

---

**From:** Rattner, Justin  
**Sent:** Monday, October 27, 2008 12:49 PM  
**To:** Hurley, Jim  
**Cc:** Prajapati, Ranna H; Schutz, Joseph  
**Subject:** RE: Status at Pixar

She was going to talk to Greg and let me know. They were speaking today as I recall. She described the situation as a "cluster" you-know-what.

If I hear from her, I will send you an update. If I don't hear from her, I will ping her again.

- JR

**From:** Hurley, Jim  
**Sent:** Monday, October 27, 2008 11:44 AM

40068DOC000020  
CONFIDENTIAL

**To:** Rattner, Justin  
**Cc:** Prajapati, Ranna H; Schutz, Joseph  
**Subject:** RE: Status at Pixar

Any update? Jim.

**From:** Rattner, Justin  
**Sent:** Thursday, October 23, 2008 6:41 AM  
**To:** Hurley, Jim  
**Cc:** Prajapati, Ranna H; Schutz, Joseph  
**Subject:** RE: Status at Pixar

Thanks, Jim, for the update. I'm at LAP report out and will speak to Renee today about this. Stay tuned.

- Justin

**From:** Hurley, Jim  
**Sent:** Wednesday, October 22, 2008 2:20 PM  
**To:** Rattner, Justin; Schutz, Joseph  
**Cc:** Prajapati, Ranna H  
**Subject:** Status at Pixar

Hi Justin,

We've run into a snag with Pixar. Remember the exchange we had about a certain Pixar employee who was mulling over the notion of working for us and finally decided to stay there? Well, apparently 3 or 4 others recently left Pixar for Renee's group, and this has Greg completely bent out of shape – he's incandescent about the possibility of their IP being used for the benefit of his competition (i.e. PDI), and by IP in this context he's not worried about their characters and stories, but the algorithms and the general wherewithal to build movie quality rendering systems.

This has become an obstacle in our discussions with Pixar, in the meeting we had with them they offered us a scene from their upcoming movie "Car Toons", we need this just to understand what we don't know we don't know about the inner workings of these kinds of engines, to see what all we need to get from them to be able to extract all the models, textures, shaders, lights etc so we can then gage how our ray tracer can compete and generate the necessary quality output. Now they are afraid that we also could potentially leak their IP to the competition.

I think the only way we're going to be able to get past this logjam is if between you and Renee we can agree that (a) we will no longer target Pixar employees for recruitment (apparently, there is evidence that we were targeting other employees explicitly, I believe Renee has stopped this already), (b) we would firewall the existing hires in SSG so they are not working with Pixar competitors (I'm not sure if they are actually working with the PDI project or not, this could just be his fear), and (c) that we would do the same – namely firewall their technology, and the people working with their technology.

What do you think? Perhaps once you and Renee are on the same page, you could give him a call to assuage his worst-case-concerns?

Jim.

PS: If you need a blow-by-blow you can read Ranna's running commentary – attached.

40068DOC000021  
CONFIDENTIAL

# EXHIBIT 9



**From:** Gelsinger, Patrick P  
**To:** Cooper, Don  
**CC:** Kennay, Jody L; Sahgal, Narendar B; Dickenson, Christina L; Johnson, Jim A; Garbus, Elliot D  
**Sent:** 10/10/2008 11:36:19 AM  
**Subject:** RE: Intel recruiting Pixar employees

Thanks Don -- this is great. No need for a complete moratorium - people who approach us are fair game. I'll reply directly and privately to GregP

Regards,  
 Pat.

-----Original Message-----

**From:** Cooper, Don  
**Sent:** Friday, October 10, 2008 10:13 AM  
**To:** Gelsinger, Patrick P  
**Cc:** Kennay, Jody L; Sahgal, Narendar B; Dickenson, Christina L; Johnson, Jim A; Garbus, Elliot D  
**Subject:** RE: Intel recruiting Pixar employees

Hi Pat,

Jason is out today so I thought I would answer your question on Pixar. We haven't mounted a strategy to specifically target Pixar employees, but I can imagine that any we came across would have been fast tracked based on their relevant skills and experience. We can take a look at any DEG hires that have come from Pixar and identify how they came to us.

In the meantime, rest assured that Staffing Mgmt has sent a note out to all of the recruiters letting them know they must STOP any direct sourcing of Pixar employees due to the sensitive nature of Intel's partnership.

Let me know if you want a complete hands-off of Pixar employees (in other words, should we avoid any who apply directly as well?). We will get back to you early next week on what we find out in regards to those people we've already hired.

I assure you we will be above reproach on this issue going forward.

Don

-----Original Message-----

**From:** Gelsinger, Patrick P  
**Sent:** Thursday, October 09, 2008 10:47 PM  
**To:** Johnson, Jim A; Cooper, Don; Garbus, Elliot D  
**Cc:** Kennay, Jody L; Sahgal, Narendar B  
**Subject:** FW: Intel recruiting Pixar employees

Please give me an update as to what is going on and proposed response.

Regards,  
 Pat.

-----Original Message-----

**From:** Greg Brandeau [mailto:brandeau@pixar.com]  
**Sent:** Thursday, October 09, 2008 5:31 PM  
**To:** Gelsinger, Patrick P  
**Cc:** Prajapati, Ranna H; Christopher Ford  
**Subject:** Intel recruiting Pixar employees

40006DOC000497  
 CONFIDENTIAL

Hi Pat,

Over the past 6-9 months, Intel has made offers to 3 engineers who work in our RenderMan team. Given that Pixar and Intel do have a close relationship, could you ask the recruiters inside of Intel to please stop recruiting our people? I am not saying the crazy thing that people at Pixar shouldn't be allowed to apply to Intel job postings, I am simply asking that you don't cold call our people. Especially Dana's team which is only 15 people in the first place!

Thanks!

-gtb

40006DOC000498  
CONFIDENTIAL

# EXHIBIT 10

---

**From:** Dickenson, Christina L  
**To:** Cooper, Don; Cunningham, Casey; Dunne, Sandra L; Jensen, Brig C; Noble, Scott A; Padilla, Anthony; Rangaswamy, Sowmya; Storm, Barbara M; Gurumurthy, Sujatha; Woolsey, Dave; Hughes, Jeanine; Knowles, Scott J; Dailey, Margaret E; Brooks, CamilleX A; Peery, TiffanyX L; Scarsellato, ValerieX; Dixon, PamX  
**Sent:** 10/10/2008 9:39:17 AM  
**Subject:** Pixar employees

Hi,  
We've been asked to avoid directly sourcing any Pixar employees... It's causing problems for DEG.  
We're still ok to talk to them if they come to us, but don't cold call or target them directly.

Thanks!


**Christina L. Dickenson**  
GAM Strategic Recruiting Manager  
Intel Corporation

40006DOC000488  
CONFIDENTIAL

# EXHIBIT 11



**DELL INTEL  
PARTNERSHIP**



November, 2007

Dell Confidential

# OUR STRATEGIC TRANSFORMATION

## FIX THE CORE

- Price optimization
- Org. structure & effectiveness
- Product development, fulfillment & ODM
- Sales productivity
- Planning processes

## RE-IGNITE GROWTH

- Consumer
- Emerging Countries
- Notebook
- Marketing

## BUILD FOR LONG-TERM

- Consumer
- Enterprise
- SMB
- Services / Solutions

Dell Confidential



# GUIDING PRINCIPLES

## **We must simplify IT for our customers by removing cost and complexity**

- Commitment to industry standards; Lowest environmental impact

## **We must grow...it's a land grab and we have to be at the forefront**

- Emerging countries; Notebook; Consumer

## **We must access all types of customers**

- Direct; Retail for Consumer and SMB; Channel for Large Commercial

## **We must re-establish our cost leadership**

- OPEX discipline; Leadership in product cost, supply chain, fulfillment

## **We must offer a broader range of products and develop for customers faster**

- Expand product lines particularly in Notebook and Enterprise
- Regional-specific and customer-targeted products
- Right touch development model and faster cycle times

## **We must invest in differentiation and innovation**

- Return to leadership in technology transitions

## **We must prioritize services and solutions**

- Support Services
- Infrastructure Consulting Services
- Deployment and Managed Services; SMB / MB Outsourcing Services

Dell Confidential





# "BIG 5" GROWTH INITIATIVES

## SMB

- Direct and channel direct
- Services and compelling products build the moat

## EMERGING COUNTRIES

- Only real full-line alternative to HP
- Focus on BRIC + next 10
- Build unique set of products
- New sales, distribution and services capabilities required

## NOTEBOOK

- Cost Leadership
- Product Leadership
- Customer Access
- GDFs

## ENTERPRISE

- Lead the industry in driving relevance in standards-based computing for virtualization, power & cooling and storage

## CONSUMER

- Product lust
- Disruptive innovation
- All in on Retail
- Annuity revenue and margin streams

Dell Confidential



# MUTUAL OPPORTUNITIES TO COLLABORATE

## SMB

- Further drive Vostro – channel and product initiatives
- Joint marketing in channel / resellers

## EMERGING COUNTRIES

- Create joint plan to address next billion PC users
  - Develop China PC, India PC
- Move beyond BRIC
- Cost and product leadership

## NOTEBOOK

- Drive wireless, mobility
- Develop low cost, emerging country products
- Innovation, leadership NBs

## ENTERPRISE

- Accelerate 1-socket tower and rack
- Drive 10GbE
- Focus on virtualization, power & cooling

## CONSUMER

- Improve communications / computing
  - Customer experience, converged product, etc
- Drive product lust – reasons to aspire to advanced features, form factor marketing, etc
- Joint marketing in retail

Dell Confidential



# MUTUAL OPERATIONAL OPPORTUNITIES

## PRODUCT & MARKETING

Align Intel technology with Dell innovative products



Enhanced product design and planning collaboration

Customer-specific products

Mutual go-to-market

Evolved IIP

## SUPPLY CHAIN

Align supply chains and fulfillment models



Inventory

Transitions

Availability / continuity

Country-specific fulfillment

## BEYOND THE BOX

Enable beyond the box revenue and margin



Content


Services

Country-specific

Dell Confidential



# EXHIBIT 12

 United States | Buy Online or Call 1-800-AWWW-DELL

[Sign In](#) | [Cart](#)



[Shop](#) [Support](#) [Community](#)

Keyword Search



[About Dell](#) [News](#) [Press Releases](#)

[PRINT](#) [SHARE](#)

## Press Releases

### Dell to Lead in Offering Dual-Core Technology

[Print](#) [Email](#) [RSS Feed](#)

Date: 2/8/2005

Round Rock, Texas

Dell, the world's leading supplier of desktop and workstation computers, today announced its intention to offer Intel's dual-core processor technology on its high-performance personal computers for consumers and businesses.

Dell has worked closely with Intel on the new technology for several years and has successfully tested dual-core processor-based systems in its engineering labs over the past few months.

"Dell is the world's preferred technology provider because we have the ability to bring the benefits of the latest technology to the everyday computer user," said John Medina, senior vice president, Dell Product Group. "Our leadership in the desktop and workstation markets demonstrates Dell's ability to deliver innovative technologies that customers value and appreciate."

With the performance gains provided by dual-core technology, extreme gamers will have the ability to play the most powerful and complex games while running multiple applications concurrently, typically without performance degradation. Multimedia enthusiasts will delight in the ability to watch a DVD while burning music and other digital content simultaneously. Workstation customers will benefit from the ability to multi-task in a complex environment and optimally run multi-threaded applications.

Dell plans to offer Dimension XPS gaming systems and Dell Precision workstations with Intel dual-core processor technology later this year.

#### About Dell

Dell Inc. (Nasdaq: DELL) is a premier provider of products and services required for customers worldwide to build their information-technology and Internet infrastructures. Company revenue for the past year totaled \$47.3 billion. Dell, through its direct business model, designs, manufactures and customizes products and services to customer requirements, and offers an extensive selection of software and peripherals. Information on Dell and its products can be obtained at [www.dell.com](http://www.dell.com).

[About Dell](#) [News](#) [Press Releases](#)

#### Shop

[Home & Home Office](#)  
[Small & Medium Business](#)  
[Large Business](#)  
[Public Sector](#)  
[Partners](#)

#### Learn

[About Dell](#)  
[Careers](#)  
[Small and Medium Business](#)  
[Solutions Center](#)

#### Support

[Support Home](#)  
[Drivers & Downloads](#)  
[Unresolved Issues](#)  
[Product Support](#)  
[Microsoft Vista Support](#)

#### Community

[IdeaStorm](#)  
[Direct2Dell Blog](#)  
[Discussion Forums](#)  
[Ratings & Reviews](#)  
[StudioDell](#)

#### My Account

[Sign-in / Register](#)  
[Order Status](#)  
[Saved Carts](#)

© 2009 Dell | [About Dell](#) | [Terms of Sale](#) | [Unresolved Issues](#) | [Privacy](#) | [About Our Ads](#) | [Contact](#) | [Site Map](#) | [Feedback](#)

# EXHIBIT 13


**news resources**
[Press Releases](#)
[Press Kits](#)
[Photos](#)
[Chip Shots](#)
[Innovation@Intel](#)
[Video / Audio](#)
[Media Center](#)
[Intel Blogs](#)
[Contact Intel PR](#)
**Corporate Information**
[About Intel](#)
[Biographies](#)
**Legal Information**
[Competition in the Innovation Economy](#)
**Search News Resources**
 
[Advanced Search >](#)

## Intel News Release

### Intel, Dell And Schlumberger To Deliver Feature Rich Platforms For The Oil And Gas Industry

SANTA CLARA, Calif., Nov. 9, 2005 – Intel Corporation today announced it is working with Dell and Schlumberger to deliver a new workstation platform specifically designed to meet the needs of the oil and gas industry.

The three companies will employ their core strengths to develop computing platforms that deliver solutions with superior data throughput, computation performance and 3D graphical performance. Specifically, Schlumberger brings to this collaborative effort an intimate knowledge of exploration and production (E&P) environments and customer needs. Dell brings the ability to design and engineer high-performance Dell Precision workstations. Intel provides the next-generation architecture that promises to deliver systems with improved performance and memory capacity. In addition, Intel® software tools are being used by Schlumberger to further optimize the new Intel technology-based solutions to deliver faster workflow and data solutions.

In today's extremely challenging E&P environment, new oil and gas reserves are increasingly difficult to find, and existing reserves are increasingly difficult to recover. To succeed, the industry must be innovative and speed up the transformation of data into decision-ready intelligence. Intel, Dell and Schlumberger are working together to deliver a high-performance computing platform for the E&P industry that meets the oil and gas industry's needs for productivity and innovation.

"Our customers are continually pushing the limits of existing workstations," said Ihab Toma, Schlumberger Information Solutions vice president of sales and marketing. "Working with Intel and Dell through their development lifecycle to deliver a new, high-performance workstation will ensure that our clients are able to realize the full potential of our latest software solutions and to make faster, more accurate, decisions."

"Together with Schlumberger, a recognized leader in the industry, Intel and Dell are able to understand the requirements dictated in a comprehensive seismic-to-simulation workflow," said John Davies, vice president, Sales and Marketing Group and general manager, Customer Solutions Group, Intel Corporation. "This improved understanding enables Intel to harness its technology and software tools to bring a game-changing solution to this industry, enabling geologists,

**Related links**
[More in this category](#)
[Corporate Information](#)
[Contact Corporate Press Relations](#)

geophysicists, and reservoir engineers to perform dynamic subsurface interpretation and modeling.”

“Dell’s goal in this collaboration is to deliver scalable, reliable and high performance workstations with standards-based technology designed to reduce complexity and help customers get more out of their IT spend,” said Neil Hand, vice president, Dell Product Group. “Our product team welcomes the challenge to advance technologies that help customers achieve new levels of performance, stability and reliability to run such demanding applications as the Schlumberger Petrel® workflow tools.”

This effort represents a deepening of the already strong technical relationship among the three companies. Today, Schlumberger routinely optimizes their software prior to release on the latest Intel® platforms to deliver faster E&P workflow and data solutions, using Intel software tools. Dell and Intel have aligned to produce advanced workstations, the Dell Precision® 470 and 670, based on the dual core technology of the Intel® Xeon® processor. Dell, the world’s leading supplier of workstations,<sup>1</sup> recently announced the availability of these Dell Precision® workstations, which enable dramatic increases in performance and power efficiency compared to single-core processors, especially when running multiple or multi-threaded applications.<sup>2</sup>

#### About Intel

Intel (NASDAQ: INTC), the world leader in silicon innovation, develops technologies, products and initiatives to continually advance how people work and live. Additional information about Intel is available at [www.intel.com/pressroom](http://www.intel.com/pressroom) and [blogs.intel.com](http://blogs.intel.com).

Intel is a trademark or registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

1. IDC WW Workstation Tracker, Q2 2005

2. Based on the SPEC CPU2000 benchmark test performed by Dell Labs in August 2005 on the Dell Precision 670 with two dual-core 2.8 Intel Xeon, 2x 2MB L2 cache processors as compared to similarly configured Dell Precision 670 with two Intel Xeon 3.6Ghz, 2x 2MB L2 cache single core processors. Actual performance will vary based on configurations, usage and manufacturing variability.

[Back to Top](#)

[Site Map](#) [RSS](#) [Jobs](#) [Investor Relations](#) [Press Room](#) [Contact Us](#)

[Terms of Use](#) [\\*Trademarks](#) [Social Media Guidelines](#) [Privacy](#) ©Intel Corporation



# EXHIBIT 14

United States | Buy Online or Call 1-800-4AWWW-DELL

Sign In | Cart



Shop Support Community

Keyword Search



About Dell News Press Releases

PRINT SHARE

## Press Releases

### Grid Starts Here with Dell, Oracle Delivering High Availability Databases for Windows Environments

**Dell Delivers Certified Configurations on Intel Platforms for Oracle Database 10g, Oracle Real Application Clusters Accelerator Services, Expedite Deployment, Customer Time-to-Value**

Print Email RSS Feed

Date: 9/21/2005  
Round Rock, Texas

Customers looking for an affordable, highly available database for Microsoft Windows Server environments now can turn to Dell for tested and validated server, software and storage configurations and comprehensive implementation services.

Dell, working with partners Oracle and Microsoft, is helping accelerate customer access to grid computing with an easy-to-use, highly available database solution designed to scale as their business grows. Mid-sized organizations and enterprises can purchase Oracle Database 10g Standard Edition and Oracle Real Application Clusters tested and validated with Microsoft Windows Server 2003 Standard Edition on a two-node Oracle cluster of Dell PowerEdge 2850 servers featuring Intel Xeon processors and Dell/EMC fibre channel storage.

"Dell, Oracle and Microsoft have combined efforts to provide customers with a highly available and scalable enterprise solution at a small-business price," said Judy Chavis, Dell's director of business development and global alliances. "These clustered solutions are easy to buy and deploy, and available from one point of contact - Dell. In addition, Dell Services can simplify, accelerate and optimize the deployment of Oracle and Windows through fixed-scope, fixed-fee offerings."

Oracle Database 10g Standard Edition and Oracle Real Application Clusters on Dell servers enables customers to run packaged or custom applications unchanged in a cluster, eliminate underused standby systems costs, reduce unplanned downtime while maximizing hardware use and stability, and improve user service and application availability levels - all key factors in accelerating customer time-to-value.

"Oracle Database 10g on Microsoft Windows and Dell PowerEdge server is another powerful offering of great value from Oracle and Dell that hits the sweet spot of affordability, reliability, and performance for many mid-sized organizations," said Oracle Vice President of Platform Alliances, Judson Althoff. "Dell and Oracle are providing the entry-level building blocks of grid computing infrastructure that will give customers an affordable platform to scale out server and storage resources as their business grows."

#### Dell: One-Stop, Comprehensive Services

Dell offers implementation services for system installation and configuration as part of the cluster's package price. Dell's services team will build each system using industry best practices and collaborations with EMC, Intel, Microsoft and Oracle. A key aspect of the solution is comprehensive management through Oracle Enterprise Manager 10g, for monitoring and managing business services and applications deployed on a grid, and EMC Navisphere management tools.

Dell will also serve as a single point of contact to address any issues relating to hardware, software and operating systems. System maintenance and monitoring is simple as Dell, Oracle, and Microsoft provide intuitive, standards-based management interfaces. Initial roll-out of the Dell Accelerator services are limited to the United States.

"Intel, Dell and Oracle are working together to deliver customers value-driven, standards-based enterprise products, solutions and accelerator services based on Intel® Xeon processor technology. The collaboration of these industry-leading companies helps to ensure customers their entire solution stack performs optimally and can scale as their business needs grow," said Neil Blecherman, director of corporate development, Customer Solutions Group, Intel Corporation. "Customers who utilize accelerator services also gain value by capitalizing on best practices from Project MegaGrid - which includes auto provisioning and scaling of compute resources with great price performance."

Configurations are priced according to processors, memory and storage and are available today worldwide. For more information on certified Oracle Database offerings for Dell PowerEdge servers, visit [www.dell.com/oracle](http://www.dell.com/oracle)

#### About Dell


Dell Inc. (NASDAQ: DELL) is a trusted and diversified information-technology supplier and partner, and sells a comprehensive portfolio of products and services directly to customers worldwide. Dell, recognized by Fortune magazine as America's most admired company and No. 3 globally, designs, builds and delivers innovative, tailored systems that provide customers with exceptional value. Company revenue for the last four quarters was \$52.8 billion. For more information about Dell and its products and services, visit [www.dell.com](http://www.dell.com).

About Dell News Press Releases

Shop	Learn	Support	Community	My Account
Home & Home Office	About Dell	Support Home	IdeaStorm	Sign-in / Register
Small & Medium Business	Careers	Drivers & Downloads	Direct2Dell Blog	Order Status
Large Business	Small and Medium Business	Unresolved Issues	Discussion Forums	Saved Carts
Public Sector	Solutions Center	Product Support	Ratings & Reviews	
Partners		Microsoft Vista Support	StudioDell	

[© 2000 Dell](#) | [About Dell](#) | [Terms of Sale](#) | [Unresolved Issues](#) | [Privacy](#) | [About Our Ads](#) | [Contact](#) | [Site Map](#) | [Feedback](#)

# EXHIBIT 15

 United States | Buy Online or Call 1-800-AWWA-DELL

Sign In | Cart

[Shop](#) [Support](#) [Community](#)

Keyword Search

[About Dell](#) [News](#) [Press Releases](#)

PRINT

SHARE

## Press Releases

### Dell Announces Energy-Saving Product Strategy

#### Energy Resource Provides Calculation Tools; Industry Challenged to Join Effort

[Print](#) [Email](#) [RSS Feed](#)

Date: 9/12/2006

New York

Dell announced its energy-efficiency product strategy today, underscoring the importance of reduced power consumption as well as enhanced performance in its product design.

The company also launched an energy resource guide at [www.dell.com/energy](http://www.dell.com/energy) to help customers protect their power needs using Dell products.

"Dell is committed to providing our customers products with the most performance per watt," said Dell CEO Kevin Rollins. "Our energy-efficient products help customers lower cost of ownership and help meet the broader goal of protecting the environment. It's critical that our entire industry address this important issue."

The focus on product energy efficiency is designed to help customers decrease electricity usage and system operating cost. Customers visiting Dell's energy resource site can view the energy-efficient features of several Dell product families and access energy calculators to help estimate power needs. A calculator for the energy-efficient Dell OptiPlex™ 745 desktop system launched today, and calculators for other product families will be available later this year. A separate datacenter capacity planner allows customers to "drag and drop" their desired product configuration into a virtual server rack and estimate power required to run their IT operations.

#### Product Energy Efficiency Features

Dell's new OptiPlex 745 features Intel® Core™ 2 Duo processors, which help deliver up to 30 percent better performance and power savings of about 40 percent versus previous generation Intel processors<sup>1</sup>. The total number of customers expected to deploy the new OptiPlex system with power-efficient Dell Energy Smart settings, an Intel Core 2 Duo processor and a Dell flat-panel monitor, are estimated to have the potential to save nearly \$1 billion every year in energy costs worldwide compared to previous-generation systems, or avoid the equivalent emissions of removing nearly one million cars from the road<sup>2</sup>. The new OptiPlex 745 also incorporates Dell's HyperCool™ thermal-management technology in the chassis to ensure better reliability and quieter, cooler operation and thus less power is required for cooling the systems during operation.

Dell's recently launched ninth-generation PowerEdge™ servers can lower system-level power consumption by up to 25 percent, which when combined with significant increases in performance, provides gains of up to 193 percent<sup>3</sup> in performance per watt. These new PowerEdge Servers incorporate several features to help reduce power consumption. Each server is built with an Intel® dual-core Xeon® E5100 series processor, which is designed to significantly increase performance while reducing power requirements, as well as high efficiency power supplies and low form factor SAS (Serial Attached SCSI) hard drives that can reduce thermal output.

All Dell Inspiron™ and Latitude™ laptops and Dell Precision™ mobile workstations now ship with sleep-state enabled, following a precedent set by the company's OptiPlex and Dell Precision desktop workstation products. This feature automatically puts the computer into a low-power "sleep" state after 15 minutes of inactivity, but allows the computer to quickly "wake-up" when needed. Sleep state enablement is a key factor of Energy Star compliance which can help reduce electricity needed to operate these systems by an estimated 70 percent.

Dell also announced that all external power supplies for its Inspiron and Latitude laptop products are designed to meet new Energy Star® requirements for computer external power supplies<sup>4</sup>.

#### About Dell

Dell Inc. (NASDAQ: DELL) listens to customers and delivers innovative technology and services they trust and value. Uniquely enabled by its direct business model, Dell sells more systems globally than any computer company, placing it No. 25 on the Fortune 500. Company revenue for the past four quarters was \$57.4 billion. For more information, visit [www.dell.com](http://www.dell.com). To get Dell news direct, visit [www.dell.com/RSS](http://www.dell.com/RSS).

[About Dell](#) [News](#) [Press Releases](#)